(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
19 April 2007 (19.04.2007)

PCT

(10) International Publication Number
WO 2007/044099 A2

(51) International Patent Classification:
G06F 15/16 (2006.01)

(21) International Application Number:
PCT/US2006/025018

(22) International Filing Date: 26 June 2006 (26.06.2006)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/166,893 24 June 2005 (24.06.2005) US

(63) Related by continuation (CON) or continuation-in-part (CIP) to earlier application:
US 11/166,893 (CON)
Filed on 24 June 2005 (24.06.2005)

(71) Applicant *(for all designated States except US)*: **AIRVANA, INC.** [US/US]; 19 Alpha Road, Chelmsford, Massachusetts 01824 (US).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **CHERIAN, Sanjay** [US/US]; 6 Maxwell Drive, Brookline, New Hampshire 03033 (US). **NG, Dennis** [US/US]; 126 Indian Meadow Drive, Northboro, Massachusetts 01532 (US). **BARABELL, Arthur J.** [US/US]; 11 Hayden Circle, Sudbury, Massachusetts 01776 (US). **RAMASWAMY, Suresh** [US/US]; 40 Old Stage Road, Chelmsford, Massachusetts 01824 (US). **GARG, Deepak** [IN/US]; 56 Stillwater Drive, Nashua, New Hampshire 03062 (US).

(74) Agents: **FEIGENBAUM, David L.** et al.; Fish & Richardson P.C., P.o. Box 1022, Minneapolis, Minnesota 55440 (US).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA,

[Continued on next page]

(54) Title: PRESERVING SESSIONS IN A WIRELESS NETWORK

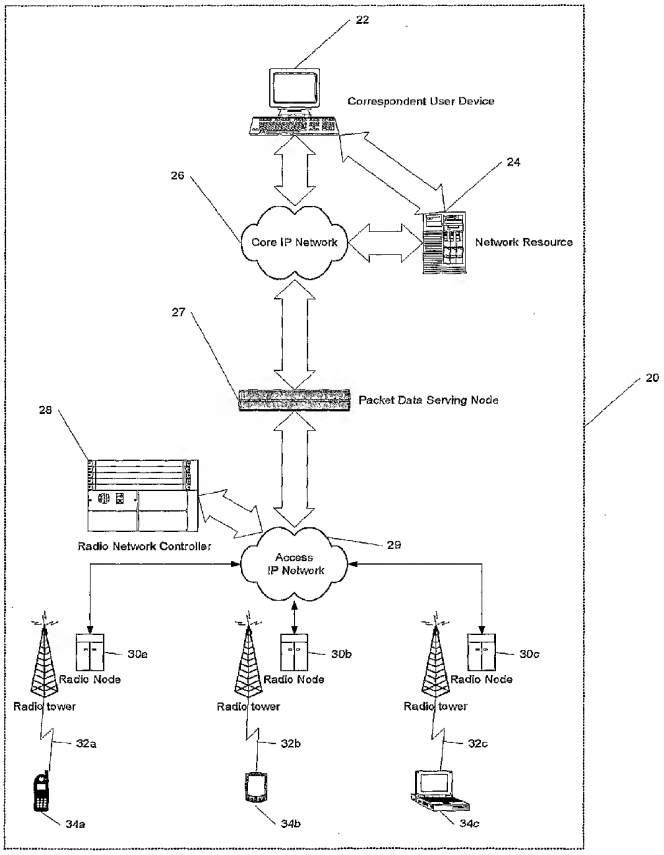

(57) Abstract: A radio network controller and methods for reestablishing sessions in a wireless network are described. At least a portion of session information associated with a first session is saved; and in response to detecting an unexpected degradation of the first session, reestablishment of the first session is triggered using the portion of the session information.

WO 2007/044099 A2

NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

— *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

— *without international search report and to be republished upon receipt of that report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

PRESERVING SESSIONS IN A WIRELESS NETWORK

TECHNICAL FIELD

This disclosure relates to the preservation of communication sessions with client devices in a wireless network communication system despite a failure in components of the communications systems.

BACKGROUND

In a fixed or mobile wireless Internet Protocol (IP) network, a client-side device establishes a session with a network-side access device to communicate with other entities in that network. The session represents the client-side device to the network, and includes information about the client-side device such as its IP address, location within a mobility area, permitted services and other such attributes required to communicate with the client-side device. The session may be partitioned into separate data elements representing IP connectivity and link-layer connectivity and these data elements may be stored on separate network elements. The session is typically present as long as the client-side device is present in the network, though other resources required for direct communication between the client-side device and the network device may only be in use when active communication is in process.

The state of direct communication between the client-side device and the network-side device is known as a connection. In cellular wireless systems, as the client device must initiate connection, a special procedure known as paging is used to allow the network to request the client device to communicate. For example, if a network device needs to send data to a client device, the network uses the information stored in a session associated with the client device to page the client device. Paging involves transmitting a message addressed to a particular terminal over a shared channel monitored by all terminals in communication with that part of the network. This page causes the client device to initiate a connection with the network, thus enabling an exchange of data. However, if a hardware or software failure of a network-side device causes the network to lose the session information (referred to as a "session breach"), the network cannot establish a direct communication with the client device, as it has lost all knowledge of the device's identity, communication parameters and location required to page that particular device. A session for which session

information has been lost due to hardware or software failure of a network-side device is referred to as a "breached session." Because the session breach is not visible at the IP (Internet Protocol) layer, the client's peers are unaware that the client is unreachable. Thus, network-initiated connection-oriented applications such as online video teleconferencing and Internet telephony are particularly vulnerable to network-side failures. Furthermore, the client device is not immediately aware that the session breach has occurred and its recovery mechanisms operate on a sufficiently long timescale that client-initiated creation of a new session cannot be counted upon to restore network reachability for that client before it is required.

SUMMARY

In one aspect, reestablishing a session may be accomplished by saving at least a portion of session information associated with a first session between an access terminal (e.g., a cellular telephone, a personal data assistant, or a laptop computer) and a first wireless network device; and in response to detecting an unexpected degradation of the first session, triggering a reestablishment of the first session using the portion of the session information.

Implementations may include one or more of the following features. Degradation may include cessation and detecting a degradation of the first session may include detecting a state (e.g., failure) of the first wireless device. Triggering a reestablishment of the first session may include transmitting, to the access terminal, a close session message. A first session may be replicated without being closed and restored upon receiving a request to open a new session from the access terminal. The triggering may comply with a 1x Evolution-Data Optimized protocol and / or be based on a load state of a second wireless network device.

Transmitting a close session message may occur immediately upon detection of a unexpected degradation of the first session and / or after receiving a request to transmit data to the access terminal. Degraded sessions may be placed in a queue and moved up in the queue in response to receiving a request to transmit data to an access terminal associated with the degraded session. Degraded sessions may include breached sessions. The session information for the session assigned to the access terminal may be deleted if the access terminal has failed to request to open a new session and / or a second wireless network device fails to reestablish the first session

after a predetermined time has elapsed. The first session may be established between the access terminal and a first wireless network device. The session information may be saved on a second wireless network device that generates and / or transmits the close session message.

A second session may be established between the access terminal and the second wireless network device; and at least a portion of the second session information that sufficient to reestablish the second session may be saved to a third wireless network device. The portion of the second session information may be sufficient to generate a close session message for the access terminal for the second session. In response to receiving a close session message, a breached session may be reestablished by closing the breached session and sending a request to open a new session.

In another aspect, a radio network controller includes a first radio node server module configured to establish a session with a first access terminal; a storage device (e.g., non-volatile random access memory, a flash memory, and a disk memory) configured to store at least a portion of the session information that is sufficient to reestablish the session; and a control mechanism configured to cause a second radio node server module device to reestablish the session with the access terminal after detecting a degradation of the session between the first radio node server module and the access terminal.

Implementations may include one or more of the following features. The session information may be sufficient to generate a close session message and /or complies with a 1x Evolution-Data Optimized protocol. The control mechanism may be configured to transmit the close session message to the access terminal and to retrieve the portion of the session information from the storage device and send the portion to the second radio node server module without causing the session to be closed. The second radio node server module may transmit a close session message immediately after the control mechanism detects a degradation of the session between the first radio node server module and the access terminal. A degradation may include a cessation. The second radio node server module may transmit a close session message only after the control mechanism receives a request to transmit data to the access terminal. Degraded sessions may be placed in a queue such that a closed session is moved to a higher entry of the queue when a request is received to transmit data to an

access terminal associated with at least one of the degraded sessions. The first radio node server module may include a first processing card and the second radio node server module may include a second processing card. The control mechanism may be implemented on a processor that is connected to the first radio node server module and to the second radio node server module through a high speed bus. The control mechanism may be implemented on the second radio node server module or on a third radio node server module.

In another aspect, reestablishing breached sessions in a wireless communications network is accomplished by placing a first session that has been breached in a queue for reestablishment of the first session; placing a second session that has been breached in the queue for reestablishment of the second session, which is prioritized below the first session in the queue; and promoting the second session above the first session in the queue in response to receiving a request to transmit data to an access terminal associated with the second session.

Implementations may include one or more of the following features. The wireless communications network may use a 1 x Evolution-Data Optimized protocol. Reestablishment of the second session may be triggered by generating and transmitting a close session message to the access terminal associated with the second session. Reestablishment may also be triggered based on a load state of a second wireless network device. The second session may be reestablished between a wireless network device and the access terminal. Triggering reestablishment of the first session may be performed after triggering reestablishment of the second session, and in some examples, only after receiving a request to transmit data to an access terminal associated with the first session. The time that the first session has spent in the queue may be monitored and the first session may be deleted if it has occupied an entry in the queue past a predetermined time period.

These general and specific aspects may be implemented using a system, a method, or a computer-readable medium, or any combination of systems, methods, and computer-readable mediums.

The details of one or more embodiments of the invention are set forth in the accompanying drawings and the description below. Other features, objects, and advantages of the invention will be apparent from the description and drawings, and from the claims.

DESCRIPTION OF DRAWINGS

FIG. 1 shows a block diagram of a wireless network system.

FIG. 2 shows internal modules of a radio network controller (RNC).

FIG. 3 depicts the RNC internal modules of FIG. 2 in greater detail.

FIG. 4 depicts the relationship between protected RNSMs and their session data.

FIG. 5A illustrates the order in which sessions are closed for proactive session closure.

FIG. 5B illustrates the order in which sessions are closed for reactive session closure.

FIG. 5C illustrates the order in which sessions are closed for combined proactive and reactive session closure.

FIG. 6 illustrates the logic of proactive session restoration performed by the Protecting RNSM of FIG. 4.

FIG. 7 illustrates the logic of reactive session restoration performed by the Protecting RNSM of FIG. 4.

FIG. 8 illustrates the logic of combined proactive and reactive session restoration performed by the Protecting RNSM of FIG. 4.

FIG. 9 depicts the Session Database of FIG. 3 and FIG. 4.

FIG. 10 illustrates the logic of proactive session closure performed by the Protecting RNSM of FIG. 4.

FIG. 11 illustrates the logic of reactive session closure performed by the Protecting RNSM of Fig. 4.

FIG. 12 illustrates the logic of combined proactive and reactive session closure performed by the Protecting RNSM of Fig. 4.

FIG. 13 illustrates the logic of session closure functions performed by the Base I/O (BIO) of FIG. 2.

DETAILED DESCRIPTION

Session breaches caused by network-side hardware and/or software failures can be remedied by saving a complete copy of each user session at an alternate location that is not likely to fail at the same time as the location hosting the primary copy of that user session. As the session data changes whenever the state of the client device changes, it must only be saved when in particular states known to be consistent between the

network and the client device. These states are unique to the wireless communication protocol in use and the internal design of the RNC. For this approach to be effective, all saved states must be such that a client device may immediately open a connection upon restoration of the session to that saved state.

Session breaches can also be remedied by transmitting a message to a client-side device instructing the client-side device to close the current session and reopen a new session with the network. For example, the 1x-Evolution Data Optimized (1xEV-DO) protocol currently defined in the IS856 family of standards, as defined by the 3GPP2 organization, provides for a "close session message". A close session message is a message that can be sent by either a client-side or network-side device to close a user session. For example, a client-side device (referred to as an "access terminal" in the 1xEV-DO protocol) may send a close session message to the network when the device is powering off or when its 1xEV-DO application is shut down. A network side device, such as a radio network controller, may send a close session message when it is shutting down or attempting to perform some form of overload control due to constrained resources. When an access terminal receives a close session message, it is required under the 1xEV-DO protocol to close its current session with the network. It will reopen a new session if it anticipates further communication with the network. Client devices that participate in network-initiated applications such as telephony always reopen new sessions as long as they are powered on. Therefore, a network-side device can use the close session message to reestablish session information lost due to a network-side failure. Moreover, because the information necessary to generate and transmit a close session message is less than the total amount of session information stored by a network-side device, the network does not need to maintain a complete backup of each session established on the network at a particular time. The benefit of this is that the network can maintain a larger number of user sessions for a given amount of memory available on the network-side device.

Referring to FIG. 1, a wireless IP network 20 includes a correspondent user device 22, a core IP network 26, a network resource 24, a radio network controller (RNC) 28, a Packet Data Serving Node (PDSN) 27, an access IP network 29, and multiple radio nodes (RN), three of which are shown 30a-30c, that each communicate with an access terminals (AT) 34a-34c using an airlink 32a-32c.

The access terminals, e.g., AT 34a, are the client-side of the wireless network and Internet protocol, and may be implemented as a mobile device such as a cellular telephone, a wireless PDA, a handheld gaming device, or a wireless laptop computer or a stationary device such as a desktop computer, a parking meter, or other fixed device with which wireless communication is desired.

The correspondent user device 22 connects to a core IP network 26 through an optional network resource 24 such as a telephony gateway and represents the device with which the ATs communicate over the wireless IP network. For example, the correspondent user device 22 may consist of an Internet telephone that permits a user at the device 22 to engage in a voice over IP call with a user associated with an AT. Similarly, if the ATs include a network of wireless-enabled parking meters, a user at the device 22 may use a graphical user interface to send messages to each parking meter instructing it to provide status information (e.g., space-filled/time paid; space-filled/time expired; space-empty/time remaining, space-empty/time expired, etc.). If the correspondent user device is a traditional telephone, it may connect through an optional network resource in the form of a media gateway and a soft switch that together form an interface between an IP network and a traditional telephone network.

The core IP network 26 is a network of devices that use the TCP/IP network protocols to exchange data. The core IP network 26 may, for example, comprise the public Internet or a private intranet. In a CDMA cellular data system, the core IP network 26 interfaces with the wireless network through a Packet Data Serving Node (PDSN) 27.

The RNC 28 communicates with the PDSN 27 and with the RNs 30a-30c over the access IP network 29 controls the radionodes' transmitters and receivers, initiates and maintains client sessions, directs data packets received from the core IP network 26, and performs other radio access and link maintenance functions such as soft handoff and sector selection. The access IP network 29 may be a private network that is designed and operated exclusively for communication between PSDN 27, RNC 28, and RNs 30a-c.

Referring to FIG. 2, the RNC 28 includes a System Controller (SC) 40, one or more Base Input/Output modules (BIO) 42, and one or more Radio Node Server Modules 46a-46c (collectively referred to as RNSMs 46) that are interconnected over a high-speed bus 44 such as a PCI, VMEbus, USB, ISA, or PXI bus or switched fabric

such as ATM (or other cell-based fabric) or Ethernet (or other packet-based fabric). Each of the three Radio Node Server Modules 46a-46c shown communicate with an AT 34a-34c using a radio node 30a-30b. In reality, each RNSM will typically use multiple radio nodes to communicate with many ATs at any particular time. For simplicity, however, only one radio node per RNSM and one AT per node are illustrated.

Each RNSM 46a, 46b, 46c is responsible for establishing and maintaining a session with the ATs that they are assigned to handle by the System Controller (SC) 40. As previously mentioned, a session must be established before data can be exchanged. During the lifetime of a session maintained by an RNSM, the RNSM continually sends session status information to the SC 40 or to an optional Session Server 38. If an RNSM fails, the SC detects the failure through heartbeat messages for which it receives no response. If an RNSM reboots or otherwise recovers, it registers its presence with the SC 40, thus alerting the SC 40 to its availability to provide service.

Referring to FIG. 3, the System Controller 40, which generally functions to perform protocol route computations, system configuration, network management, and centralized signaling, also maintains a session lookup database 70 of the current sessions established between the RNSMs 46a-46c and the ATs. This database 70 is normally not a backup of all session information maintained by the RNSMs but includes information sufficient to identify the RNSM on which a particular session is hosted. The session may be located by multiple identifiers, such as International Mobile Subscriber Identity (IMSI), Unicast Access Terminal Identifier (UATI), or hardware ID. This database 70 may also be implemented on a separate Session Server 38. The use of a Session Server 38 enables location of sessions across multiple RNC chassis. The session lookup database 70 is accessed and updated by a session lookup database application 71.

The Base I/O Module 42 functions to connect the RNC 28 with the access IP network 29 and the PDSN 27. It receives packets destined for the PDSN 27 and the RNs 30 from the RNSMs 46 and SC 40 and transmits them over its network interfaces to the access IP network 29. It also receives packets from the PDSN 36 and the RNs 30 from the access IP network 29 and routes them to the RNSMs 46 and SC 40. For packets received from the PDSN 36, the BIO 42 extracts the PDSN/PCF-specific identifier (PSI) field from the packet and looks it up in its PSI forwarding table 62 to find the identity of the RNSM (e.g., RNSM 46a) on which the corresponding session is

located. For packets received from the RNs 30 over the airlink access channel, the BIO 42 extracts the Unicast Access Terminal Identifier (UATI) from the packet and looks it up in its UATI forwarding table 63 to find the identity of the RNSM (e.g., RNSM 46a) on which the corresponding session is located.

The RNSM (e.g., RNSM 46a) functions to perform the wireless-specific protocol functions necessary to implement the RNC's call processing and data handling capabilities. The call processing functions are implemented by the Call Control 72 component. The Call Control 72 component interacts with the Packet Control Function (PCF) Signaling 64 component that implements the portions of call processing that pertain to the interface with the PDSN 27. The Call Control 72 component and the PCF Signaling 64 component maintain the state of the sessions that they set up as part of call processing in the Active Session Database 66a.

To preserve user reachability after session breach, a Session Backup Manager 68a on RNSM 46a interacts with a Session Backup Manager 68b on RNSM 46b to maintain a Backup Session Database 67a on RSNM 46a that is a backup of the Active Session Database 66b on RNSM 46b and a Backup Session Database 67b on RNSM 46b that is a backup of the Active Session Database 66a on RNSM 46a. This diagram shows a specific case of two RNSMs 46a and 46b protecting each other.

Protection Models

Referring to FIG. 4, the RNSMs 46 are configured such that RNSM 46a maintains a Backup Session Database 67a to protect the Active Session Database 66b on RNSM 46b and RNSM 46b maintains Backup Session Database 67b to protect the Active Session Database 66a on RNSM 46a. RNSM 46c and RNSM 46d have a similar configuration where each RNSM maintains a backup of the other's session datatbase. In some embodiments, SC 40 maintains a Backup Session Database for one or more of the RNSMs 46.

In the event of failure of RNSM 46a, RNSM 46b distributes its Backup Session Database 67b to itself and the other RNSMs 46c and 46d as shown by the gray shaded extensions to the Active Session Databases 66b, 66c and 66d. The options for processing these sessions are described more below. Note that, in this protection strategy, an even number of RNSMs 46 are populated into the RNC 28 to guarantee that every RNSM 's session database is protected. If a new RNSM is added to the

RNC 28, it may be paired up with an odd remainder RNSM that was existing in the RNC or can remain unprotected until an additional RNSM is added to the RNC.

In some implementations, each RNSM in the RNC protects one other RNSM, for example, the RNSM located to its right in the RNC's chassis with the rightmost RNSM in the chassis is responsible for protecting the leftmost RNSM in the chassis. In this manner, an RNSM protection ring is established. Thus, in this protection scheme any number of RNSMs (greater than or equal to 2) can be populated in an RNC with protection. As new RNSMs are added to the RNC, they are inserted into the protection ring, such that two other RNSMs break their existing protection relationship and establish a protection relationship with the newly added RNSM.

In some other implementations, each RNSM in the RNC may partially protect multiple other RNSMs in the RNC such that each RNSM is fully protected by a set of other RNSMs.

The RNSM population of the RNC chassis can be determined at power-on of the RNC, even though the RNSMs may complete their software initialization in a random order. However, because the RSNM population can be determined at power-on, in some implementations, a protection arrangement is selected at power-on but the protection does not begin until after a configurable hold-down time to allow all the RNSMs to initialize. In the event of a late appearance of an RNSM, e.g., due to failure at the RNC initialization time, it is treated as an insertion of a new RNSM module into the RNC chassis.

Recovery Models

The traditional approach to implementing any form of redundancy-based reliability consists of maintaining a backup of the data to be protected and restoring that data upon detection of failure of the primary copy of the data. FIG. 5A shows the traditional service order for restoring this data, referred to as Proactive Service Order. In this approach, the backup data is ordered in some fashion and is restored according to its existing order. It is prior art in wireline networking and I believe in wireless networking as well. Restoration by sorting UATI may be unique so perhaps we should take that sentence out.

If the restoration of a session takes a significant amount of time, the proactive approach may result in longer outages of service than tolerable as the user sessions are

restored in an order independent of which users are actually active and would have benefited from session restoration. In a CDMA system, network-side reachability is required only when the network, in the form of the packet data servicing node (PDSN), has a packet to send to the AT. Thus, it is possible to defer restoration of any given session until a packet destined to the AT represented by that session is received from the PDSN 27. This is referred to as Reactive Session Restoration and its service order is depicted in FIG. 5B. This approach minimizes the outage perceived by an average session but has a constant value for the outage perceived by a session that was required to be restored. It also has the property that session state has to be preserved for a long time after failure of an RNSM while waiting for packets to arrive from the PDSN 27 prior to the expiration of the session lifetime.

A hybrid scheme is proposed, referred to as Proactive plus Reactive. In this approach, sessions are restored according to their order in the Backup Session Database as in the Proactive Approach. If, however, a packet arrives from the PDSN 27 for a session that has not already been closed, that session is promoted to the head of the list for restoration. FIG. 5C shows this service order.

Session Backup and Restoration

One general approach to preserving user reachability in the event of session breach is to save all information associated with user sessions to an alternate location and restore that information in the event of an RNSM failure. This general approach can be implemented in many different ways, some of which are described below.

For example, in a first approach, all sessions hosted on a particular RNSM, known as the primary RNSM, may be stored on another RNSM, known as the backup RNSM, that is dedicated to the function of storing the session data for the primary RNSM. In this approach, each RNSM in the RNC chassis has a designated backup RNSM that is idle until it takes over the function of a failed primary RNSM. The backup RNSM constantly monitors the correct functioning of the primary RNSM and, in the event of failure of the primary RNSM, assumes the function of that primary RNSM and activates its copy of the user sessions that existed on the primary RNSM. The processing capacity of the RNC is, therefore, constrained to that of half the number of RNSMs in that RNC. This approach is known as 1:1 RNSM Redundancy.

In a second approach, a group of primary RNSMs is protected by a single backup RNSM that stores a copy of the session state for all of those primary RNSMs simultaneously. This backup RNSM constantly monitors the correct functioning of all its primary RNSMs and, in the event of failure of any primary RNSM, assumes the function of that primary RNSM and activates its copy of the user sessions that existed on that primary RNSM. The processing capacity of the RNC is, therefore, reduced by one RNSM for every group of protected primary RNSMs. This approach is known as 1:N RNSM Redundancy, where N is the number of primary RNSMs protected by a single backup RNSM.

A third approach involves storing a copy of the session state for any given RNSM on a single RNSM that is itself a primary RNSM. In this model, pairs of RNSMs are associated so that each is the backup RNSM for the other. This model is known as the "buddy system" and is depicted in FIG. 4 and described previously. RNSMs may be associated symmetrically, where one RNSM protects another RNSM and that RNSM protects the first RNSM, or asymmetrically, where the RNSM protecting a particular RNSM may be protected by a different RNSM.

A fourth approach involves storing a copy of the session state for any given RNSM across all other RNSMs in the RNC. Thus, every RNSM is maintaining its own sessions and a copy of some of the sessions from every other RNSM in the RNC. In the event of an RNSM failure, each of the other RNSMs in the RNC activates its copy of the user sessions associated with the faulty RNSM.

A fifth approach involves storing a copy of the session state for all RNSMs in the RNC on a centralized resource within the RNC, for example the system controller.

A sixth approach involves storing a copy of the session state for all RNSMs in the RNC on a dedicated session server 38 external to the RNC.

Regardless of the particular backup mode, session state may be stored in several ways. First, it may be stored to Random Access Memory (RAM). Second, it may be stored to non-volatile Random Access Memory in the form of battery-backed up RAM or FLASH memory. Finally, it may be stored to a hard disk drive.

FIG. 6 shows a process for proactive session restoration performed by a backup RNSM (e.g., RNSM 46a). The backup RNSM continually issues (304) heartbeat messages to the active RNSM (e.g., RNSM 46b) while performing (302) normal call processing functions. If the active RNSM is working properly, it sends an

acknowledgement of the heartbeat messages to the backup RNSM. If the backup RNSM receives (306) an acknowledgement from the active RNSM, it continues to perform (302) normal call processing functions. If the backup RNSM does not receive (306) an acknowledgement from the active RNSM, the active RNSM is faulty and its sessions are breached. Thus, the backup RNSM begins the session recovery process. During this process, the backup RNSM triggers (308) the rehoming of all RNs that were served by the failed active RNSM to working RNSMs and starts a timer. The timer may be used to terminate the session restoration process in situations where session restoration may not terminate normally. The backup RNSM then retrieves the sessions from its backup session database and restores (314) them evenly to each of the working RNSMs in the RNC, including itself. In some embodiments, the backup RNSM randomly allocates individual sessions or groups of sessions to itself and the other RNSMs. For example, if there were 8 RNSMs in an RNC and 20000 sessions per RNSM, failure of one RNSM may result in the restoration of 2857 sessions to each RNSM on average. If a particular working RNSM indicates that it is overloaded, the backup RNSM may not assign additional sessions to that RNSM, in order to avoid exacerbating the overload condition. As each session or group of sessions is restored, the backup RNSM updates (316) the UATI and PSI forwarding tables of the BIOs so that subsequent traffic destined for that session will be routed properly. Upon determining (310) that either the timer has expired or that all of the breached sessions have been restored, the backup RNSM terminates (312) the restoration process, which includes deleting all unrecovered sessions. The backup RNSM also enters a non-protecting mode, in which it no longer sends heartbeat messages to other RNSMs.

FIG. 7 shows a process for reactive session restoration performed by a backup RNSM, as shown in FIG. 4. The protecting RNSM 46a executes the performing (302), issuing (304), determining (306), and restoring (308) processes described for the proactive restoration process 300 illustrated in FIG. 6.The backup RNSM then updates (332) the BIO's UATI and PSI forwarding tables so that messages intended for the faulty RNSM are routed to the backup RNSM. In some embodiments, the BIO's UATI and PSI forwarding tables are updated using a single message by previously configuring the BIO to know the backup RNSM for every active RNSM. The backup RNSM then waits for a packet with a UATI or PSI for a breached session to be received while continually checking (310) for timer expiration or completion of the session

recovery process. Upon receipt (334) of a packet with such a UATI or PSI, the backup RNSM retrieves the corresponding session from its backup session database and restores (314) it to one of the working RNSMs in the RNC, including itself. In some embodiments, the backup RNSM randomly allocates individual sessions or groups of sessions to itself and the other RNSMs. If a particular working RNSM indicates that it is overloaded, the backup RNSM may not assign additional sessions to that RNSM, in order to avoid exacerbating the overload condition. As each session or group of sessions is restored, the backup RNSM updates (316) the UATI and PSI forwarding tables of the BIOs so that subsequent traffic destined for that session will be routed to the correct RNSM for that session. Upon determining (310) that either the timer has expired or that all of the breached sessions have been restored, the backup RNSM terminates (312) the restoration process, which includes deleting all unrecovered sessions. The backup RNSM also enters a non-protecting mode, where it is no longer sending heartbeat messages to other RNSMs.

FIG. 8 shows a process for combined reactive and proactive session restoration performed by a backup RNSM. The protecting RNSM 46a executes the performing (302), issuing (304), determining (306), and restoring (308) processes described for the proactive and reactive closure processes 300 and 330 illustrated in FIGS. 6 and 7. The backup RNSM then updates (332) the BIO's UATI and PSI forwarding tables so that messages intended for the faulty RNSM are routed to the backup RNSM. The backup RNSM determines if a UATI or PSI for a breached session has been received while continually checking (310) for timer expiration or completion of the session recovery process. Upon receipt (334) of a packet with such a UATI or PSI, the backup RNSM retrieves the sessions from its backup session database and restores them evenly to each of the working RNSMs in the RNC, including itself. If no such packet has been received, the backup RNSM retrieves (314) the next breached session or group of sessions from its backup session database and restores them evenly to each of the working RNSMs in the RNC, including itself. In some embodiments, the backup RNSM randomly allocates individual sessions or groups of sessions to itself and the other RNSMs. If a particular working RNSM indicates that it is overloaded, the backup RNSM may not assign additional sessions to that RNSM, in order to avoid exacerbating the overload condition. As each session or group of sessions is restored, the backup RNSM updates (316) the UATI and PSI forwarding tables of the BIOs so

that subsequent traffic destined for that session will be routed properly. Upon determining (310) that either the timer has expired or that all of the breached sessions have been restored, the backup RNSM terminates (312) the restoration process, which includes deleting all unrecovered sessions. The backup RNSM also enters a non-protecting mode, where it is no longer sending heartbeat messages to other RNSMs. In this manner, session restoration proceeds in an orderly fashion, with an acceleration to the head of the processing queue of any sessions for which immediate service is requested by the PDSN 27 or an RN (e.g., RN 30a).

Session Closure

As described earlier, the 1xEV-DO protocol also permits repair of session breach by the closure of a breached session from the network side, resulting in an automatic re-establishment of the session by the AT.

As will be explained in more detail below, for session closure, the session database is preferably not a backup of all session information maintained by the respective RNSMs, but only includes information sufficient to generate and transmit close session messages to the ATs in the event of a RNSM failure. With reference to FIG. 2, if a protecting RNSM 46a receives an indication that a protected RNSM 46b may have failed, the protecting RNSM 46a will direct other RNSMs (e.g., RNSM 46c) to transmit close session messages to the effected ATs. For example, if the protecting RNSM 46a receives an indication that RNSM 46b has failed, it will use the information stored in the backup session database 67a to generate a close session message for AT 34b and assigns another RNSM (e.g., RNSM 46c) to transmit the close session message to AT 34b. When AT 34b receives the close session message, it will terminate its session with RNSM 46b and open a new session with one of the operational RNSMs.

The Base I/O (BIO) 42 receives data packets from the PDSN 27 and routes them to the appropriate RNSM for delivery to an AT. Upon notification of an RNSM failure, the protecting RNSM 46a directs the BIO 42 to flag the faulty RNSM to prevent the BIO 42 from sending any further data packets to the faulty RNSM when an application tries to contact a breached session. If the BIO 42 receives a data packet from the PDSN 27 that is addressed to a faulty RNSM 46b, the BIO 42 forwards the data packet to the RNSM 46a that was protecting the faulty RNSM 46b. The protecting

RNSM 46a preferably saves the data packet or a subset of information contained within the data packet until a new session for the AT is established on another RNSM or until the new-session establishment procedure times out. Normally the protecting RNSM 46a will attempt to reestablish lost sessions according to their order in a queue; however if the protecting RNSM 46a receives an indication from the BIO 42 that an application is trying to contact a particular session in the queue, the protecting RNSM 46a will promote that session to the top of the queue and service that session immediately, as described for the combined proactive and reactive service order 170 shown in FIG. 5C. In this manner, the protecting RNSM 46a uses notification sent from the BIO 42 to prioritize the close session messages.

Referring to FIG. 3, the protecting RNSM 46a maintains a session database 66a that contains information for all sessions served by the faulty RNSM 46b. In this particular implementation, the session database 66a is accessed by protection-related RNSM processes via the Session Backup Manager 68a software module.

As shown in FIG. 9, the session database 66a stores the following information for each session 252:

(1) a Unicast Access Terminal Identifier (UATI) 254, which is an identification code that uniquely identifies an AT 34 when it is registered with the access network;

(2) a Hardware Identifier (HwID) 256, which is a unique and permanent identification code of the physical hardware of the AT 34;

(3) a PDSN/PCF Specific Identifier (PSI) 258, which is a unique identification assigned to each AT 34 that has established a session with the network to allow the PCF subsystem of the RNC and the PDSN to identify to each other the particular session referenced by a data transmission;

(4) a control channel cycle 260 identifying the periodic interval at which the AT monitors broadcast transmissions from the network when it does not have an active radio resource uniquely assigned to it;

(5) a set of AT sectors 262 that maps each AT 34 to the particular sectors that the AT was monitoring when it last sent a route update to the network 46; and

(6) Flags 264 that indicate the occurrence of an RNSM failure.

To generate and transmit a close session message to an AT under the current 1xEV-DO protocol, the RNC maintains at least the AT's UATI, PSI, control channel cycle and last sectors received as part of a route update. The UATI is used to identify a

particular AT when sending the close session message over the control channel. The Hardware Identifier is used to identify an AT in the session database so that the RNC can detect if a session no longer requires closure as it has already been closed and replaced with a new session by an AT-initiated connection attempt. The PSI identifies the session to close when a packet is received from the PDSN 27 destined to an AT whose session has been breached. The control channel cycle is used to transmit a close session message to the AT. The AT sectors are required to optimize the session closure process. If a large number of sessions are to be closed but the location of each AT is not known, the close session message for each AT must be sent over all sectors controlled by the RNC that contains the faulty RNSM. This consumes significant amounts of processing capacity and control channel bandwidth. Knowledge of which sectors from which the AT last reported pilots allows targeting of the close session message to a smaller number of sectors for each AT. The information stored in the session database 66 is sufficient to reestablish lost sessions resulting from an RNSM failure.

Referring again to FIG. 3, the RNSM, e.g., RNSM 46a, includes a call control process 72 that monitors the status of each AT being served by the RNSM. As the status of ATs change, the call control 72 updates the Active Session Database 66a. The updates trigger the Session Backup Manager 68 to replicate the whole or part of the recently updated session onto the backup RNSM if the change to the Active Session Database results in the session entering a "syncable" state that is different from the last state that was updated to the backup RNSM. For example, if a new session is created or deleted by an AT, the call control process 72 creates or deletes a corresponding session in the session database 66a. The RNC 28 exchanges data with the core IP network 26 through an Radio-Packet (R-P) interface that carries user traffic. A single user traffic connection within this interface is an A10 as defined in the 1xEV-DO Standard. If an A10 interface is created or deleted for a session, the call control 72 updates the session database 66a which triggers the Session Backup Manager 68 to update the Backup Session Database 67 on the backup RNSM. If an AT 34 sends a route update including radio sectors that are different from the set of sectors which it included in its last route update, or changes the control channel cycle, the call control 72 updates the session database 66a. As the Session Backup Manager 68 detects update messages from the call control processes, it updates the backup session database

67 accordingly. The Session Backup Manager 68 also maintains the timers required for managing the session close and retry processes.

An RNSM communicates with the PDSN 27 through the Packet Control Function (PCF) 64. The PCF 64 controls the transmission of packets between the RNC 28 and the PDSN 27. The PCF 64 interfaces to the PDSN 27 through an A10 interface that carries user traffic between the PCF 64 and the PDSN 27. The RNC 28 opens an A10 for each session created for an AT 34. Upon detection of an RNSM failure, the PCF is configured to instruct the PDSN 27 to stop accounting for the activity of the A10s belonging to the faulty RNSM 46 but leave the effected A10s intact. This instruction can be accomplished by transmitting an ActiveStop message 76 to the PDSN 27. This step functions to maintain accurate user billing information in the core IP network as the ActiveStop command is a notification that billable user activity has ceased. Because the PCF 64 is not closing the A10s, they will be kept alive indefinitely if no other action is taken. Therefore, the A10s are closed only after a session is closed or when the entire session closure process has timed out. They are reopened when a new session is created by an AT-initiated connection attempt. The Session Backup Manager 68 is also configured to notify the BIO 42 of the affected A10s by, for example, sending the BIO 42 a flag notification to label the A10 entries belonging to the faulty RNSM (e.g., RNSM 46a).

The BIO 42 maintains a PSI table 62 that maps ATs to RNSMs and thus allows the BIO to route data packets bound for an AT to the appropriate RNSM. A PSI table 62 is a lookup structure indexed by PSI that returns the identity of the RNSM where the session addressed by that PSI is located.

The BIO also includes a Fast Path component 60 that facilitates generation of close session messaging upon detection of an RNSM failure. The fast path component receives notification of faulty RNSMs, and, in response, updates the PSI forwarding table 62, where the affected sessions are flagged. The fast path component 60 also sends received data packets to the Session Backup Manager 68 of the protecting RNSM 46a when it receives a packet from the core network labeled with a PSI whose corresponding RNSM is flagged for session closure in the PSI table. The Session Backup Manager 68 issues the GenerateSessionClose message, updated with the information from the Session Database 80 to a chosen RNSM 46 which then signals the AT 34 to close the session. When the BIO 42 receives a subsequent data packet 78

destined to a faulty RNSM 46, the BIO will discard the packet 78 if it has already initiated session closure.

When an RNSM fails, the sessions supported by the faulty RNSM are entered into a queue. The RNC can process close session messages for the sessions waiting in the queue in several ways. For example, as shown in FIG. 5A, the RNC 28 executes close session messaging using a proactive process 110, in which breached sessions are serviced immediately upon awareness of an RNSM failure according to their order in the queue 200. If a user's session is closed and reestablished before any applications attempt to reach the session, the user will never notice that his session was breached. However, if an application attempts to reach a breached session, that application must wait until all preceding queued sessions have been serviced before that particular session is reestablished. For example, suppose an application attempts to contact the fifth session in the queue 200. The application must wait for the preceding sessions to be processed before the fifth session can be reestablished, regardless of whether or not the preceding sessions require immediate service. If the fifth session is not reestablished before the application times out, the application will give up and terminate. Thus if many sessions are waiting to be closed, some users may find that an application tried to contact their sessions but could not because their sessions were at the end of a queue of sessions to be closed.

In a reactive session closure process, the RNC 28 reduces the frequency of application timeouts by transmitting close session messages to only those ATs that an application tries to reach. An example of a reactive session closure process 150 is seen in FIG. 5B. The fifth session in the queue 200 is highlighted as urgent 204 to indicate that an application is attempting to reach that session. All other queued sessions are marked as non-urgent 202. The fifth session is serviced immediately while each of the remaining queued sessions must wait until an applications attempts contact. If after a predetermined time, a queued session is not contacted by an application, the RNC times that session out and deletes it from the queue.

Because the reactive closure process prioritizes urgent sessions above non-urgent sessions, application timeouts are less likely to occur during reactive session closure than during proactive session closure. Furthermore, reactive session closure conserves processing resources that would have been required to close non-urgent sessions. On the other hand, reactive session closure causes all users to experience a

brief period of unreachability between the time at which network-initiated traffic arrived for them and the time at which it could be delivered.

A third method of session closure integrates the proactive and reactive processes to form a combined proactive and reactive process. For example, as shown in FIG. 5C, a combined proactive and reactive process 170 services non-urgent sessions 202 according their order in the queue 200. If a session becomes urgent 204, the process 170 promotes 208 that session to the top of the queue 200 for immediate servicing. The advantage of the combined scheme 170 is that the closure process can be paced so that it does not significantly disrupt normal system operation yet can be facilitated immediately if session closure becomes urgent.

FIG. 10 shows a process 370 for performing proactive session closure. The protecting RNSM 46a continually issues (304) heartbeat messages to the active RNSM 46b while performing normal call processing functions. If the active RNSM 46b is working properly, it sends an acknowledgement of the heartbeat messages to the protecting RNSM 46a. If the protecting RNSM 46a receives (306) an acknowledgement from the active RNSM 46b, it continues to perform (302) normal call processing functions. If the protecting RNSM 46a does not receive an acknowledgement from the active RNSM 46b, the active RNSM 46b is the active RNSM is faulty and its sessions are breached. Thus, the protecting RNSM 46a begins the session closure process. Immediately after an RNSM failure, the RNSM 46b reboots. Upon reboot, the RNSM 46b notifies the SC 40 that it has booted, allowing the SC 40 to begin using it for new sessions.

Upon the indication of an RNSM failure, the protecting RNSM 46a triggers the rehoming 308 of all RNs that were served by the faulty RNSM 46b. The RNSM 46a sends (372) a notification of failure of RNSM 46b to the BIO 42 to initiate a session closure process described in FIG. 13. Upon receiving the notification of failure, the BIO flags all sessions being served by the faulty RNSM 46b. The protecting RNSM 46a also notifies (372) the SC 40 that the RNSM 46b has failed. Upon notification, the SC 40 updates its Session Lookup Database 70 to reflect that the session's existence is now known on the protecting RNSM 46a. Prior to the failure of RNSM 46b, the protecting RNSM 46a had been continually storing a subset of the session information for all sessions in the RNC as given to it by the Session Backup Manager 68a. The session information is later used to reconstitute the breached session. The PCF 64

sends an ActiveStop command 76 to the PDSN 27 (374) on behalf of each breached session and starts a timer (374). The timer may be used to terminate the session closure process in situations where session closure may not terminate normally. The ActiveStop command 76 instructs the PDSN 27 to stop billing the A10s for the sessions that were active but to leave those same A10s intact. When a session is to be closed, the protecting RNSM 46a assigns a flagged session to a working RNSM (e.g., RNSM 46c) and sends the session information to the working RNSM. The protecting RNSM 46a sends the session information to a selected working RNSM (314) instructing the working RNSM to issue a session closure command (i.e., a close session message) to the AT 34. If the AT 34 requests a new session before the timeout occurs, the RNC 28 assigns the AT 34 to a working RNSM thus reestablishing the breached session. When the working RNSM registers the new session with the SessionDB Lookup App 71 on the SC 40, the Session DB Lookup App 71 determines that the session is a duplicate because it has the same HardwareID as an existing session. The Session DB Lookup App 71 responds to the working RNSM with an acknowledgement of the registered session and notification of the RNSM on which the duplicate resides – in this case, the protecting RNSM 46a. At this time, the working RNSM sends a message to the protecting RNSM 46a indicating that the session has been closed and identifying the session by its hardware ID. The protecting RNSM 46a then deletes that session from its list of sessions awaiting closure. However, if the AT 34 does not request a new session before the timeout, the protecting RNSM 46a deletes the timer and all saved session information for sessions that have not already been closed. The protecting RNSM 46a determines (378) whether all breached sessions, served by the faulty RNSM 46b, have been reassigned or whether the entire session closure timeout has expired. If the determination (378) is negative, the protecting RNSM 46a retrieves (314) the next session from its backup session database. In some embodiments, the backup RNSM randomly allocates individual sessions or groups of sessions to itself and the other RNSMs. If a particular working RNSM indicates that it is overloaded, the backup RNSM may not assign additional sessions to that RNSM, in order to avoid exacerbating the overload condition. As each session or group of sessions is closed, the backup RNSM updates (316) the UATI and PSI forwarding tables of the BIOs so that subsequent traffic destined for the session will be routed to the working RNSM to which the session is assigned for closure. Upon determining (378) that either the timer

has expired or that all of the breached sessions have been closed, the protecting RNSM 46a terminates (376) the closure process, which includes deleting the timer and all saved session information for sessions that have not already been closed. The protecting RNSM 46a enters a non-protecting mode, in which it no longer sends heartbeat messages to other RNSMs.

To reduce the impact of the recovery process on normal service delivery, a reactive session closure process 390 is implemented. FIG. 11 illustrates a reactive session closure process 390 performed by the protecting RNSM 46a. The protecting RNSM executes the performing (302), issuing (304), determining (306), notifying (372), and sending (374) processes described for the proactive closure process 370 illustrated in FIG. 10.

The protecting RNSM 46a sends the flagged session information to the BIO 42 to initiate a session closure process described in FIG. 13. Instead of initiating session closure immediately for every flagged session as was done in the proactive scheme 370, the protecting RNSM 46a waits for an application to attempt contact with the session from the core network side before initiating a session closure. The protecting RNSM 46a first starts a timer to limit the duration of time during which it will attempt to close sessions reactively (374). The BIO 42 notifies the protecting RNSM 46a that an application is attempting contact with a particular breached session by sending the protecting RNSM 46a a RecoverSession message that identifies the breached session. The RNSM 46a determines (392) whether a RecoverSession message has been received. If no RecoverSession message is received, the protecting RNSM 46a resumes its process of waiting (392) for RecoverSession messages from the BIO 42 until a timeout occurs After receiving the RecoverSession message, the protecting RNSM 46a assigns the session identified by the BIO 42 to a working RNSM. The protecting RNSM 46a retrieves the session identified by the BIO 42 from the Backup Session DB 67a and sends (336) the saved session information to the working RNSM indicating that the session is to be closed. The working RNSM sends a close session message to the AT 34. If the AT 34 requests a new session before the timeout occurs, the RNC 28 assigns the AT 34 to a working RNSM thus reestablishing the breached session. When the working RNSM registers the new session with the SessionDB Lookup App 71 on the SC 40, the Session DB Lookup App 71 determines that the session is a duplicate because it has the same HardwareID as an existing session. The

Session DB Lookup App 71 responds to the working RNSM with an acknowledgement of the registered session and notification of the RNSM on which the duplicate resides – in this case, the protecting RNSM 46a. At this time, the working RNSM sends a message to the protecting RNSM 46a indicating that the session has been closed and identifying the session by its hardware ID. The protecting RNSM then deletes that session from its list of sessions awaiting closure. However, if the AT 34 does not request a new session before the timeout, the protecting RNSM 46a deletes the timer and the saved session information. The protecting RNSM 46a determines (378) whether all breached sessions, served by the faulty RNSM 46b, have been reassigned or whether the entire session closure timeout has expired. If the determination (378) is negative, the protecting RNSM 46a waits for another RecoverSession message 80 from the BIO 42 (392) before initiating another session closure. The protecting RNSM 46a retrieves the session identified by the BIO 42 from the Backup Session DB 67a and sends (336) the saved session information to the working RNSM indicating that it is to be closed and starts a timer. The working RNSM sends a close session message to the AT 34. As each session or group of sessions is closed, the protecting RNSM 46a updates (316) the UATI and PSI forwarding tables of the BIOs so that subsequent traffic destined for the session will be routed to the working RNSM to which the session is assigned for closure. Upon determining (378) that either the timer has expired or that all of the breached sessions have been closed, the protecting RNSM 46a terminates (376) the closure process, which includes deleting the timer and all saved session information for sessions that have not already been closed. The protecting RNSM 46a enters a non-protecting mode, in which it no longer sends heartbeat messages to other RNSMs.

FIG. 12 illustrates a combined proactive and reactive session closure scheme 400 that integrates the proactive 370 and the reactive 390 session closure schemes described in FIG. 10 and FIG. 11, respectively. Both the reactive session closure process 390 of FIG. 11 and the combined closure process 400 of FIG. 12 use the RecoverSession message sent from the BIO 42 to prioritize close session messaging. The protecting RNSM 46a executes the performing (302), issuing (304), determining (306), notifying (372), and sending (374) processes described for the proactive and reactive closure process 370 and 390 illustrated in FIGS. 10 and 11.

In the combined process 400, the protecting RNSM 46a determines (378) whether all breached sessions in queue 200 have been reassigned or whether the entire session closure timeout has expired. If the determination (378) is negative, the protecting RNSM 46a determines (392) whether a RecoverSession message 80 from the BIO 42 has been received.

If the protecting RNSM 46a receives a RecoverSession message from the BIO 42, the protecting RNSM 46a will immediately process (336) the session identified by the BIO 42, thus effectively promoting that session to the front of the queue 200. If the protecting RNSM 46a does not receive a RecoverSession message from the BIO 42, the protecting RNSM 46a will process (314) the sessions according to their order in the queue 200 at a user-configurable rate. The rate may be selected to optimize a trade-off between the speed at which sessions are closed and the processing power required to close the sessions.

As each session or group of sessions is closed, the protecting RNSM 46a updates (316) the UATI and PSI forwarding tables of the BIOs so that subsequent traffic destined for the session will be routed to the working RNSM to which the session is assigned for closure. Upon determining (378) that either the timer has expired or that all of the breached sessions have been closed, the protecting RNSM 46a terminates (376) the closure process, which includes deleting the timer and all saved session information for sessions that have not already been closed. The protecting RNSM 46a enters a non-protecting mode, in which it no longer sends heartbeat messages to other RNSMs.

Referring to FIG. 13, the close session messaging process 420 performed by the BIO 42 is shown. While performing (422) normal procedures, the BIO 42 continually determines (424)whether a RNSM failure notification has been received from a protecting RNSM 46a. The BIO 42 flags (426) all session entries in its PSI table 62 and its UATI table 63 for the faulty RNSM 46b and updates the location of the sessions in the session entries to point to the protecting RNSM 46a. The BIO 42 receives data packets 78 from the PDSN 27 and determines (428) whether the packets are directed toward a faulty RNSM. If the packets are directed toward a faulty RNSM, the BIO 42 determines (434) whether the session, to which the packets belong, are flagged as "pending recovery." If the session is not flagged, the BIO 42 flags (440) the session as "pending recovery," starts a timer, and sends a RecoverSession message to the

protecting RNSM 46a. The RecoverSession message notifies the protecting RNSM 46a that an application has attempted to contact a breached session. The BIO 42 determines (436) whether the timer has expired. Subsequent packets received for the flagged sessions while the timer is not yet expired are discarded (438). If the timer has expired, a new RecoverSession is sent (440) to the protecting RNSM 46a and the timer is restarted. If the BIO 42 determines (428) that a received packet is not directed toward a faulty RNSM, the BIO 42 determines (430) if the timers for any other breached sessions that are pending recovery have expired. If the timers have not expired, the BIO 42 continues receiving packets and determining (428) whether they are directed toward faulty RNSMs. If a timer for a session has expired, the BIO 42 restarts the timer and generates (432) a new RecoverSession message for the session.

The RNC 28 could be implemented by a number of different hardware configurations. One possible configuration is a dedicated chassis containing multiple processing cards, where each card performs the functions of either the SC 40, the BIO 42, or the RNSM (e.g., RNSM 46a). In a preferred embodiment, the processing card is a Compact-PCI or Advanced TeleCommunications Architecture (A-TCA) Card containing non-volatile RAM, flash memory, and a flash disk. A second configuration is a blade server. A blade server consists of a chassis containing multiple cards where each card is equivalent to a complete single-board computer, except that common resources such as hard disk drives and power supplies are shared between all the cards. The RNC 28 could also be implemented on standalone servers arranged as a collection of independent computers in which each computer plays the role of an RNC 28. In this standalone-server configuration, each computer consolidates the functions of the SC 40, the BIO 42, and the RNSM (e.g., RNSM 46a) into a single processor. An alternate embodiment of the standalone server model utilizes each computer as either an SC, RNSM or BIO. Another configuration for the RNC 28 could be an integrated system containing embedded processors that perform the functions of the SC 40, the BIO 42, or the RNSM (e.g., RNSM 46a). A further configuration is an application-specific integrated circuit (ASIC) in which all functions of the RNC 28 are performed by a single chip. The PCF 64 could also be implemented as a virtual PCF and reside on all RNSMs 46 simultaneously or on a given RNSM (e.g., RNSM 46a) with the ability to be relocated to another RNSM(e.g., RNSM 46b).

Other embodiments are within the scope of the following claims. For example, the steps described in FIG. 6, FIG. 7, FIG. 8, FIG. 10, FIG. 11, FIG. 12 and FIG. 13 could be performed in an order that is different than the ordering shown in the figures.

WHAT IS CLAIMED IS:

1. A method comprising:

saving at least a portion of session information associated with a first session between an access terminal and a first wireless network device;

in response to detecting an unexpected degradation of the first session, triggering a reestablishment of the first session using the portion of the session information.

2. The method of claim 1 wherein degradation comprises cessation.

3. The method of claim 1 wherein triggering a reestablishment of the first session further comprises:

transmitting to the access terminal a close session message.

4. The method of claim 1 further comprising replicating the first session without closing the first session.

5. The method of claim 1 wherein detecting a degradation of the first session comprises detecting a state of the first wireless device.

6. The method of claim 5 wherein the state comprises failure.

7. The method of claim 4 further comprising:

restoring the first session upon receiving a request to open a new session from the access terminal.

8. The method of claim 1 wherein the triggering complies with a 1x Evolution-Data Optimized protocol.

9. The method of claim 3 wherein transmitting a close session message occurs immediately upon detection of a unexpected degradation of the first session.

10. The method of claim 3 wherein transmitting a close session message occurs after receiving a request to transmit data to the access terminal.

11. The method of claim 1 wherein the reestablishment of the first session is triggered based on a load state of a second wireless network device.

12. The method of claim 3 further comprising:

placing degraded sessions in a queue for transmitting a close session message; and

moving a queued degraded session up in the queue in response to receiving a request to transmit data to an access terminal associated with the degraded session.

13. The method of claim 12 wherein the degraded sessions comprise breached sessions.

14. The method of claim 1 wherein the access terminal comprises at least one of: a cellular telephone, a personal data assistant, or a laptop computer.

15. The method of claim 3 further comprising deleting the session information for the session assigned to the access terminal if the access terminal has failed to request to open a new session after a predetermined time has elapsed after transmitting the close session message to the access terminal.

16. The method of claim 4 further comprising deleting the session information for the session assigned to the access terminal if a second wireless network device fails to reestablish the first session with the access terminal after a predetermined time has elapsed after sending the portion of the session information.

17. The method of claim 1 further comprising:

establishing the first session between the access terminal and a first wireless network device.

18. The method of claim 3 wherein the session information is saved on a second wireless network device.

19. The method of claim 18 wherein the close session message is generated by the second wireless network device.

20. The method of claim 19 wherein the close session message is transmitted by the second wireless network device.

21. The method of claim 20 further comprising:

establishing a second session between the access terminal and the second wireless network device; and

saving at least a portion of the second session information to a third wireless network device, wherein the portion of the second session information is sufficient to reestablish the second session between the access terminal and the third wireless network device.

22. The method of claim 20 wherein the portion of the second session information is sufficient to generate a close session message for the access terminal for the second session.

23. A method comprising reestablishing a breached session in response to receiving a close session message.

24. The method of claim 23 wherein reestablishing the breached session further comprises:

closing the breached session; and

sending a request to open a new session.

25. A radio network controller comprising:

a first radio node server module configured to establish a session with a first access terminal;

a storage device configured to store at least a portion of the session information that is sufficient to reestablish the session; and

a control mechanism configured to cause a second radio node server module device to reestablish the session with the access terminal after detecting a degradation of the session between the first radio node server module and the access terminal.

26. The radio network controller of claim 25 wherein degradation comprises termination.

27. The radio network controller of claim 25 wherein the session information is sufficient to generate a close session message and the control mechanism is further configured to transmit the close session message to the access terminal.

28. The radio network controller of claim 25 wherein the control mechanism is further configured to retrieve the portion of the session information from the storage device and send the portion to the second radio node server module without causing the session to be closed.

29. The radio network controller of claim 25 wherein the control mechanism is configured to comply with a 1x Evolution-Data Optimized protocol.

30. The radio network controller of claim 27 wherein the second radio node server module transmits a close session message immediately after the control mechanism detects a degradation of the session between the first radio node server module and the access terminal.

31. The radio network controller of claim 30 wherein degradation comprises cessation.

32. The radio network controller of claim 27 wherein the second radio node server module transmits a close session message only after the control mechanism receives a request to transmit data to the access terminal.

33. The radio network controller of claim 27 further comprising a queue for transmitting a close session message wherein degraded sessions are placed, the queue moving a closed session to a higher entry in response to receiving a request to transmit data to an access terminal associated with at least one of the degraded sessions.

34. The radio network controller of claim 33 wherein the degraded sessions comprise breached sessions.

35. The radio network controller of claim 25 wherein the first radio node server module comprises a first processing card and the second radio node server module comprises a second processing card.

36. The radio network controller of claim 25 wherein the storage device comprises at least one of a non-volatile random access memory, a flash memory, and a disk memory.

37. The radio network controller of claim 25 wherein the control mechanism is implemented on a processor, the processor connecting to the first radio node server module and the second radio node server module through a high speed bus.

38. The radio network controller of claim 27 wherein the control mechanism is implemented on the second radio node server module.

39. The radio network controller of claim 27 wherein the control mechanism is implemented on a third radio node server module.

40. The radio network controller of claim 28 wherein the control mechanism is implemented on the second radio node server module.

41. The radio network controller of claim 28 wherein the control mechanism is implemented on a third radio node server module.

42. A method for reestablishing breached sessions in a wireless communications network, the method comprising:

placing a first session that has been breached in a queue for reestablishment of the first session;

placing a second session that has been breached in the queue for reestablishment of the second session, wherein the second session is prioritized below the first session in the queue; and

promoting the second session above the first session in the queue in response to receiving a request to transmit data to an access terminal associated with the second session.

43. The method of claim 42 wherein the wireless communications network uses a 1 x Evolution-Data Optimized protocol.

44. The method of claim 42 further comprising triggering a reestablishment of the second session.

45. The method of claim 44 wherein triggering a reestablishment comprises generating and transmitting a close session message to the access terminal associated with the second session.

46. The method of claim 42 further comprising reestablishing the second session between a wireless network device and the access terminal.

47. The method of claim 45 wherein reestablishment is triggered based on a load state of a second wireless network device.

48. The method of claim 45 further comprising:

triggering reestablishment of the first session after triggering reestablishment of the second session.

49. The method of claim 45 further comprising:

triggering reestablishment of the first session after triggering reestablishment of the second session only after receiving a request to transmit data to an access terminal associated with the first session.

50. The method of claim 45 further comprising:

monitoring the time that the first session has spent in the queue; and

deleting the first session if it has occupied an entry in the queue past a predetermined time period.

51. A computer readable medium having instructions stored thereon, that, when executed by a processor, cause the processor to:

save information associated with a first session with a wireless access terminal on a wireless network; and;

in response to detecting an unexpected degradation of the first session, triggering a reestablishment of the first session using the saved information.

52. The computer-readable medium of claim 51 wherein triggering a reestablishment comprises transmitting a close session message to the access terminal, the close session message instructing the access terminal to open a new session.

53. The computer-readable medium of claim 51 having further instructions that cause the processor to restore the first session upon receiving a request to open a new session from the access terminal.

54. The computer-readable medium of claim 51 having further instructions that cause the processor to replicate the first session without closing the first session.

55. The computer-readable medium of claim 51 wherein the wireless network uses the 1 x Evolution-Data Optimized protocol.

56. The computer-readable medium of claim 51 wherein the reestablishment of the first session is triggered based on a load .

57. The computer-readable medium of claim 51 further causing the processor to prioritize closed network sessions in a wireless communications network, the processor being caused to:

place a first session that has been breached in a queue for reestablishment of the first session;

place a second session that has been breached in the queue for reestablishment of the second session, wherein the second session is prioritized below the first session in the queue; and

promote the second session above the first session in the queue in response to receiving a request to transmit data to an access terminal associated with the second session.

58. The method of claim 57 further comprising:

triggering reestablishment of the first session after triggering reestablishment of the second session.

59. The method of claim 57 further comprising:

triggering reestablishment of the first session after triggering reestablishment of the second session only after receiving a request to transmit data to an access terminal associated with the first session.

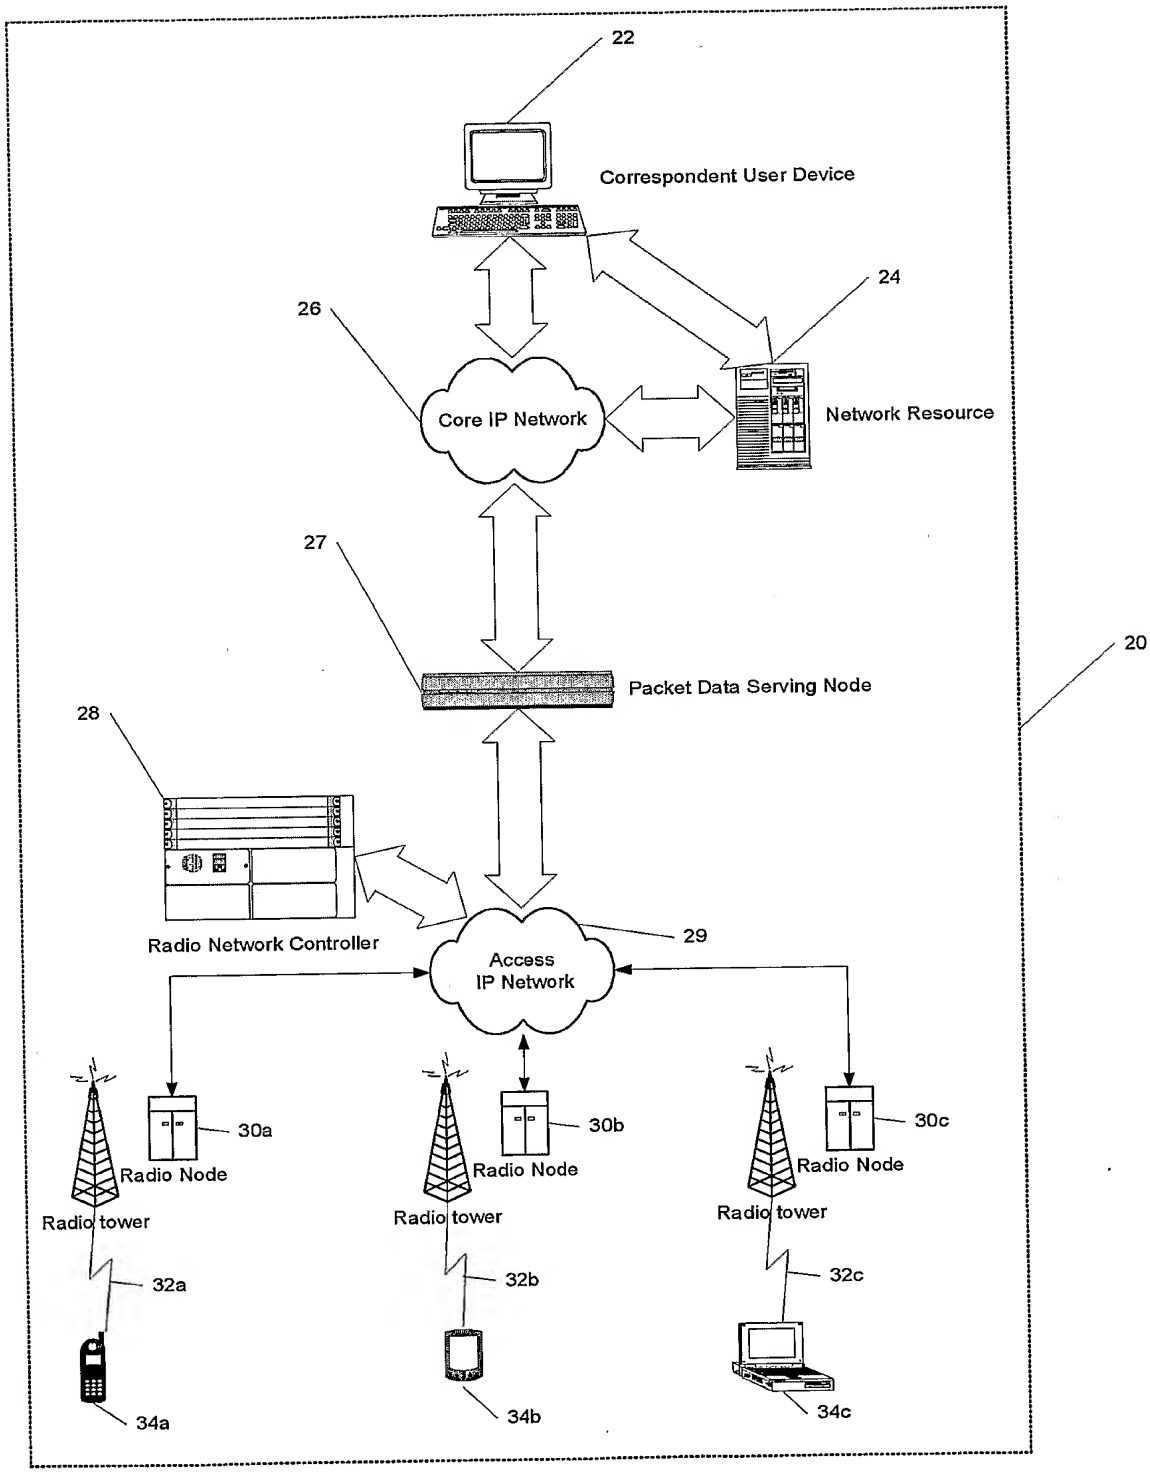
22
Correspondent User Device
24
26
Core IP Network
Network Resource
27
Packet Data Serving Node
20
28
Radio Network Controller
29
Access
IP Network
30a
Radio Node
Radio tower
32a
34a
30b
Radio Node
Radio tower
32b
34b
30c
Radio Node
Radio tower
32c
34c

FIG. 1

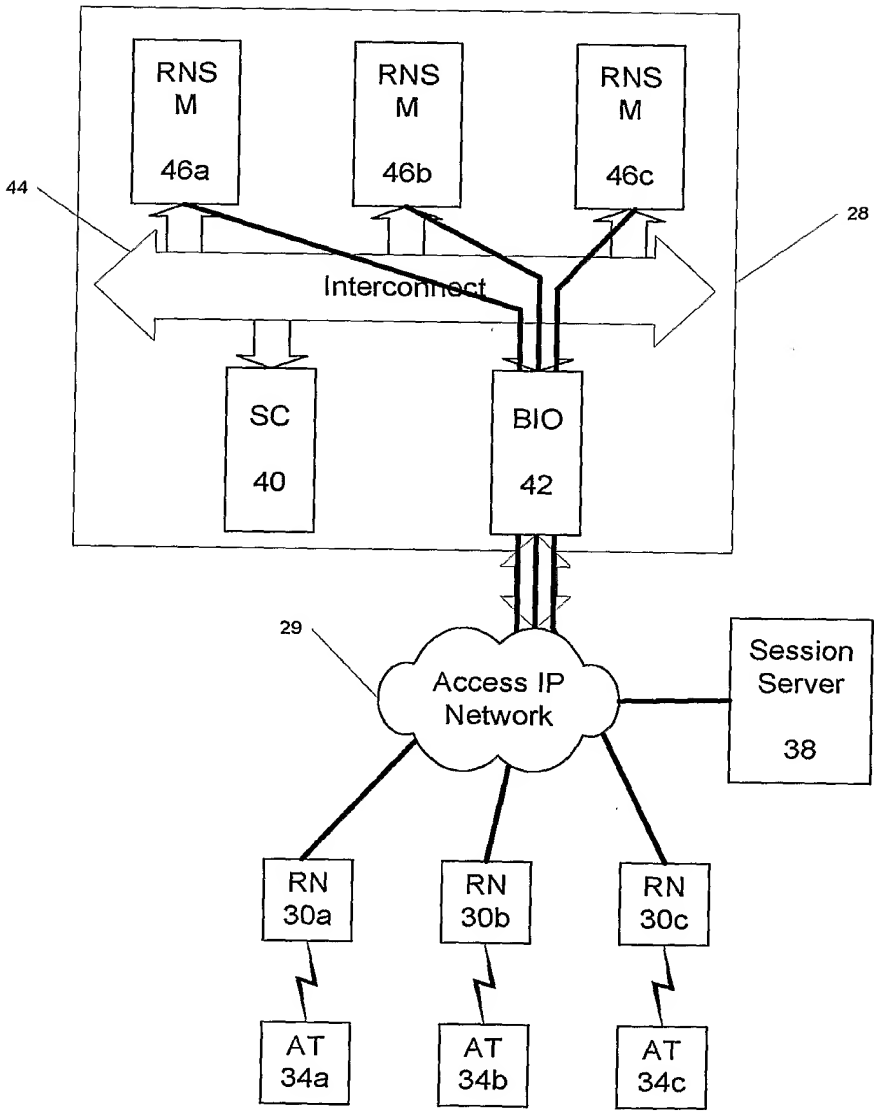
RNS M 46a
RNS M 46b
RNS M 46c
44
28
Interconnect
SC 40
BIO 42
29
Access IP Network
Session Server 38
RN 30a
RN 30b
RN 30c
AT 34a
AT 34b
AT 34c

FIG. 2

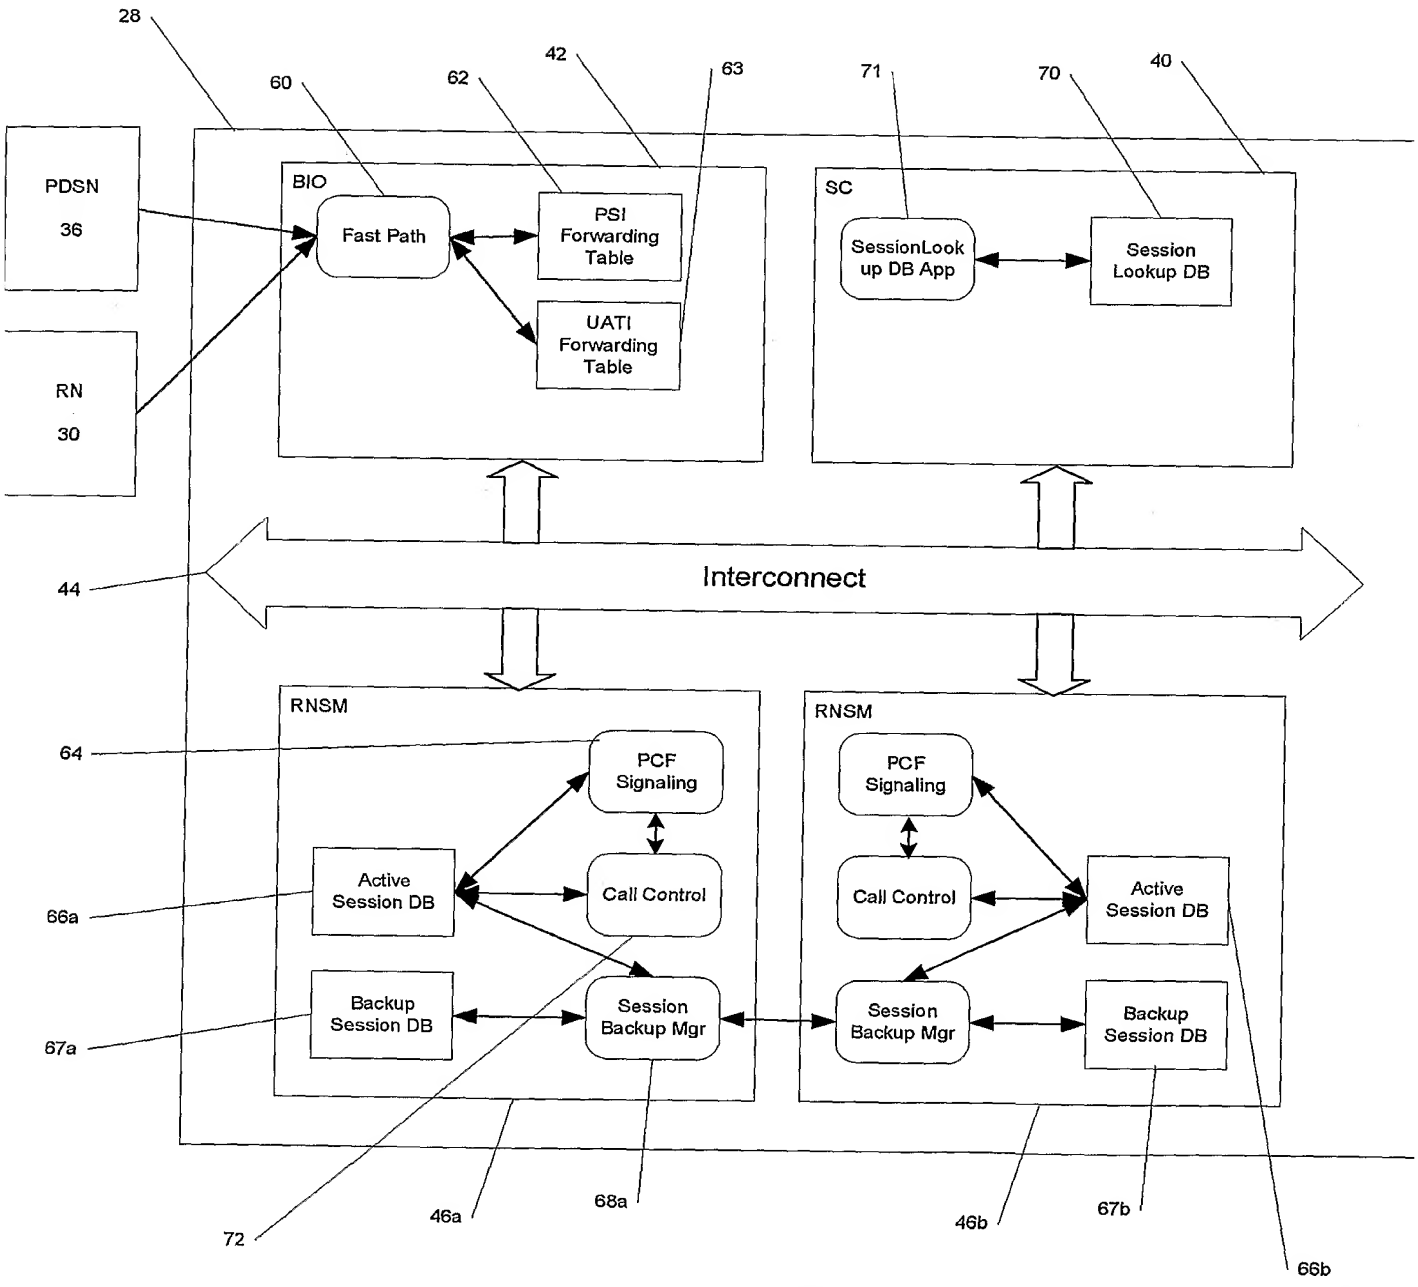
28
60
62
42
63
71
70
40
PDSN
36
RN
30
BIO
Fast Path
PSI Forwarding Table
UATI Forwarding Table
SC
SessionLook up DB App
Session Lookup DB
44
Interconnect
RNSM
64
PCF Signaling
Active Session DB
66a
Call Control
Backup Session DB
67a
Session Backup Mgr
RNSM
PCF Signaling
Call Control
Active Session DB
Session Backup Mgr
Backup Session DB
72
46a
68a
46b
67b
66b

FIG. 3

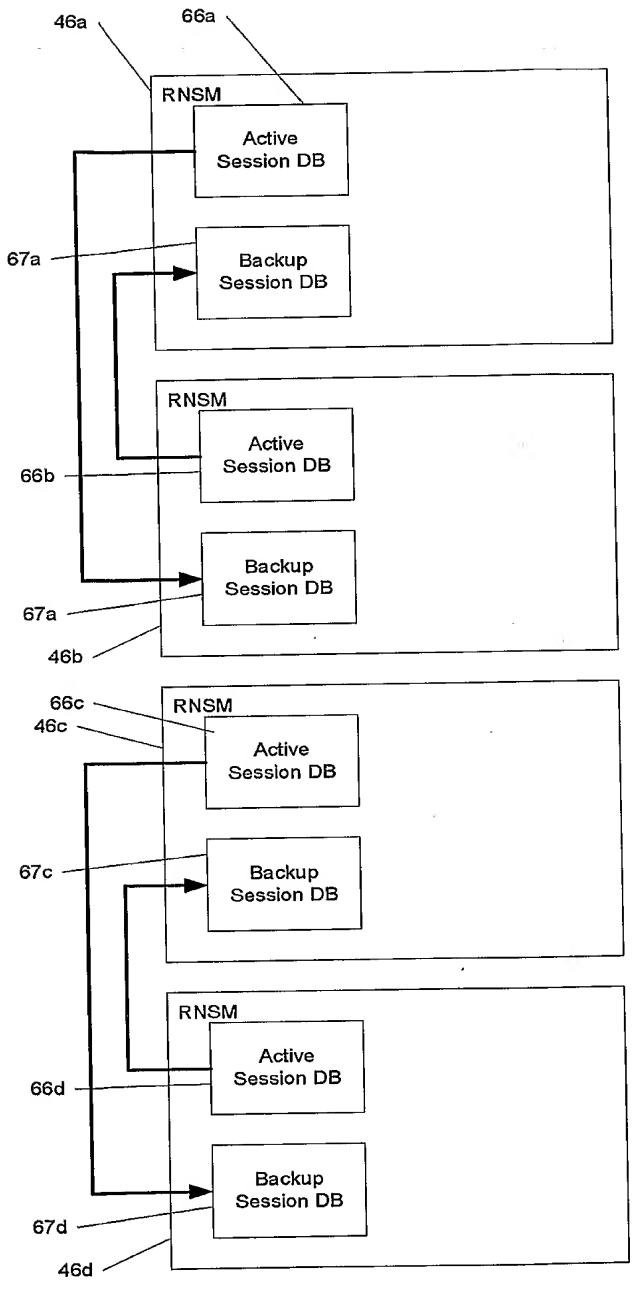

Before Failure

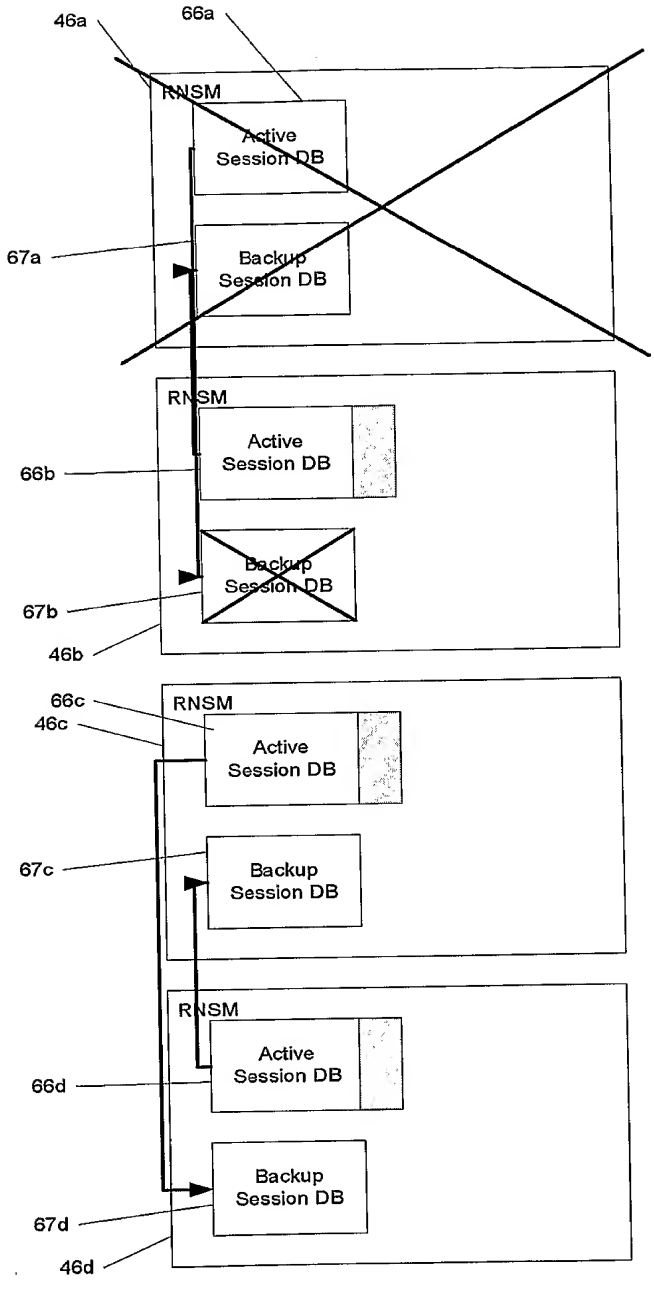

After Failure

FIG. 4

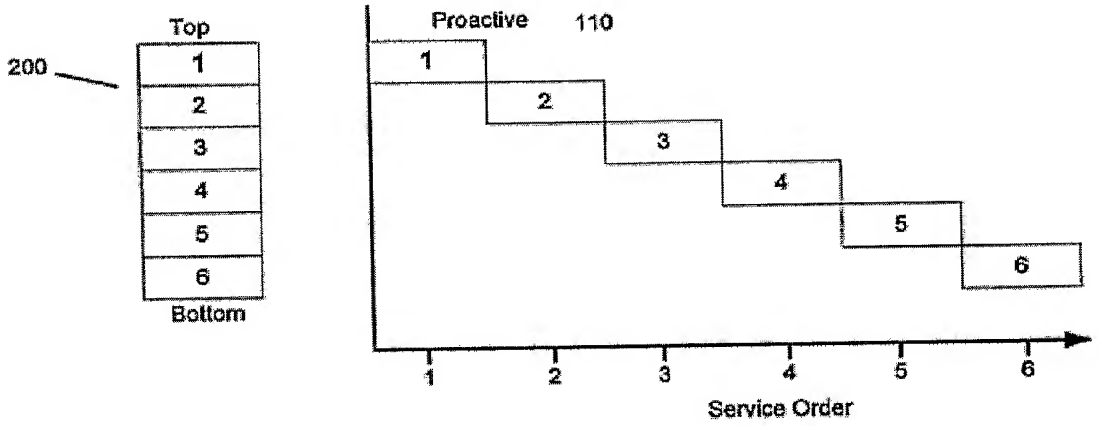
Top
200
1
2
3
4
5
6
Bottom
Proactive
110
1
2
3
4
5
6
1
2
3
4
5
6
Service Order

FIG. 5A

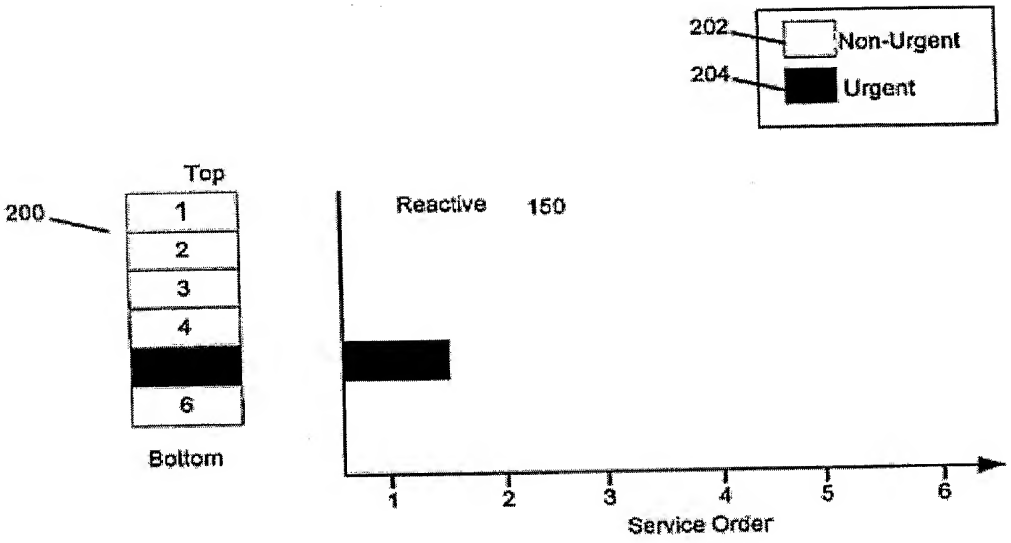
202
Non-Urgent
204
Urgent
Top
200
1
2
3
4
6
Bottom
Reactive
150
1
2
3
4
5
6
Service Order

FIG. 5B

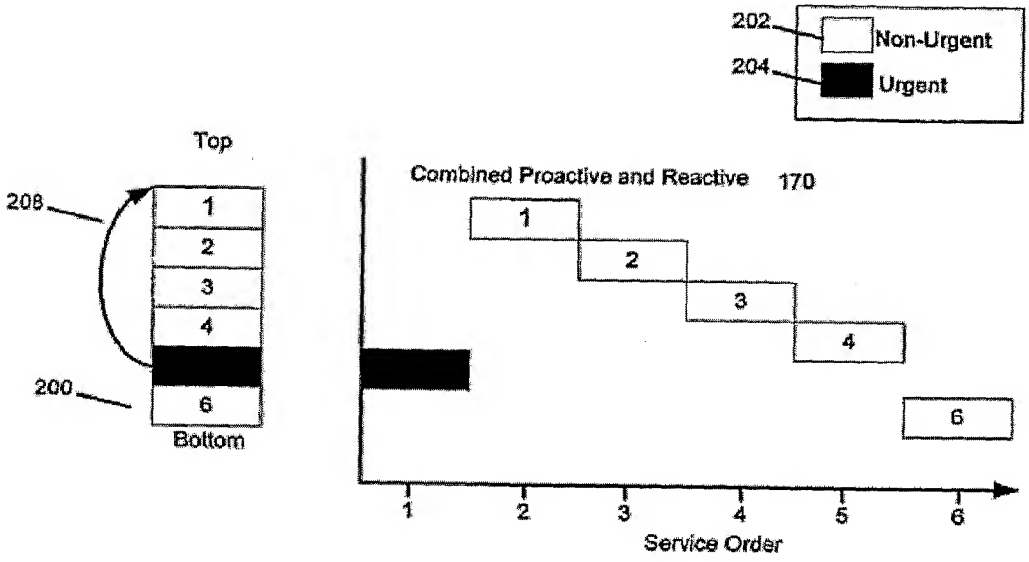
202
Non-Urgent
204
Urgent
Top
208
1
2
3
4
200
6
Bottom
Combined Proactive and Reactive 170
1
2
3
4
6
1
2
3
4
5
6
Service Order

FIG. 5C

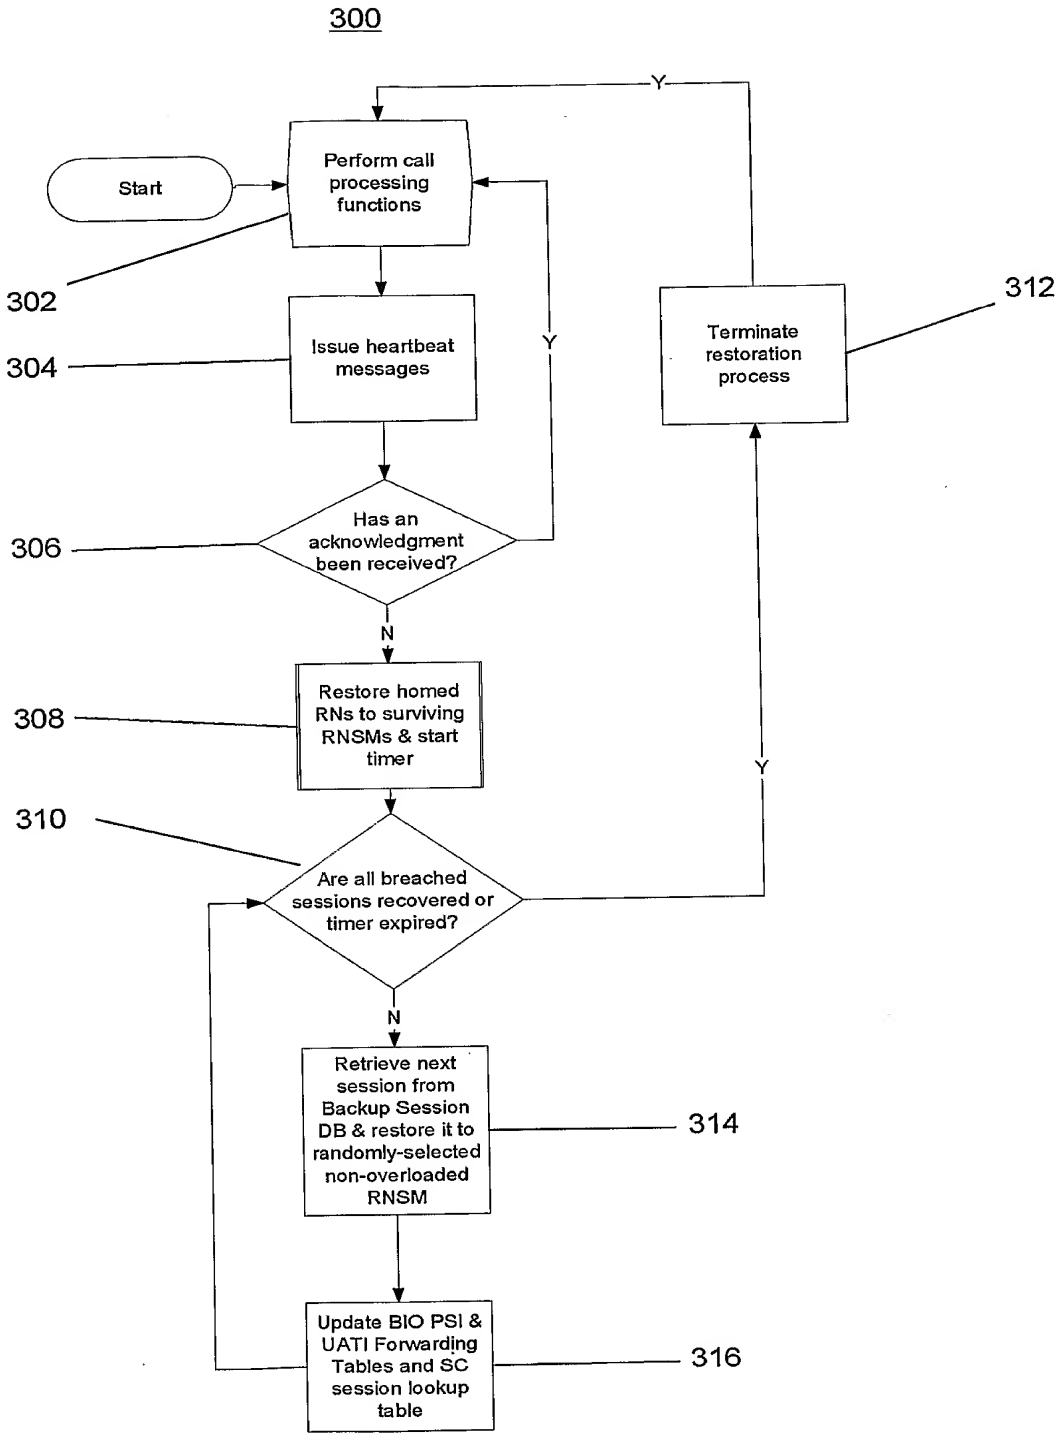
300
Start
Perform call processing functions
302
Issue heartbeat messages
304
Has an acknowledgment been received?
306
Y
N
Restore homed RNs to surviving RNSMs & start timer
308
Are all breached sessions recovered or timer expired?
310
Y
N
Retrieve next session from Backup Session DB & restore it to randomly-selected non-overloaded RNSM
314
Update BIO PSI & UATI Forwarding Tables and SC session lookup table
316
Terminate restoration process
312

FIG. 6

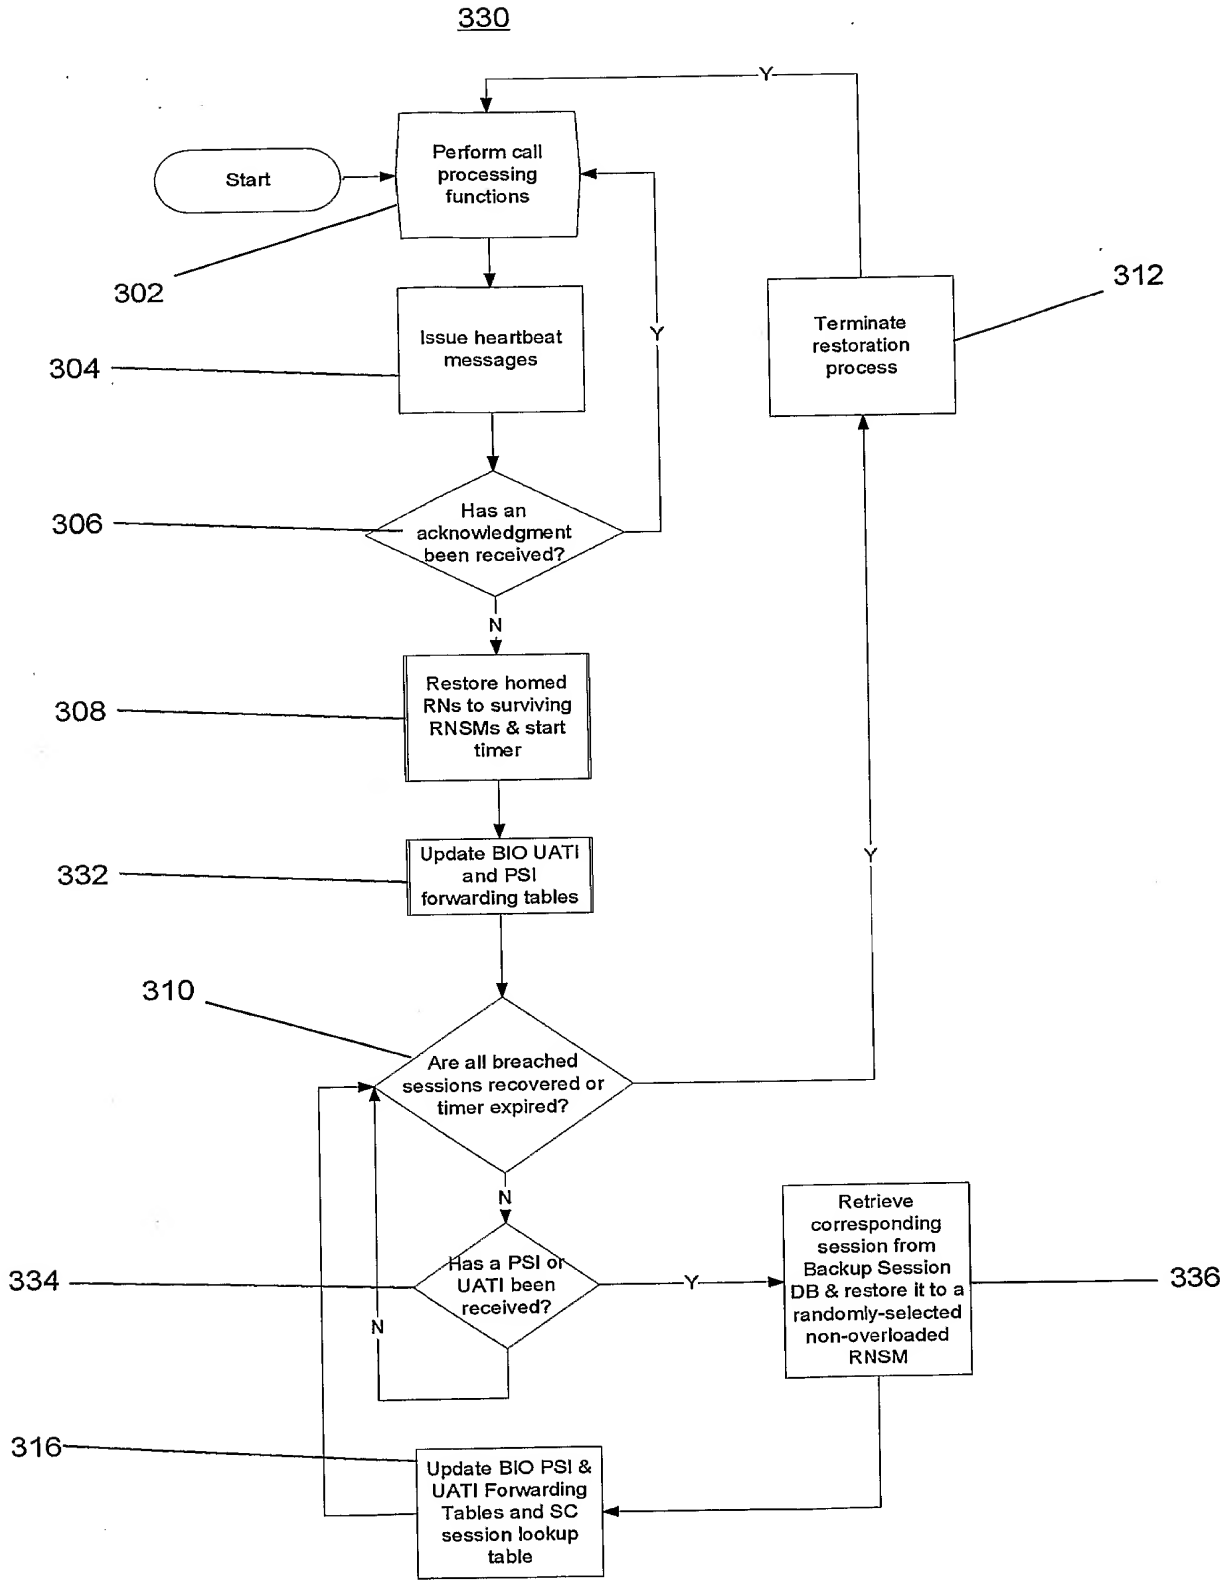
330
Start
Perform call processing functions
302
Issue heartbeat messages
304
Has an acknowledgment been received?
306
Y
N
Restore homed RNs to surviving RNSMs & start timer
308
Update BIO UATI and PSI forwarding tables
332
310
Are all breached sessions recovered or timer expired?
N
Y
Has a PSI or UATI been received?
334
N
Y
Retrieve corresponding session from Backup Session DB & restore it to a randomly-selected non-overloaded RNSM
336
Terminate restoration process
312
316
Update BIO PSI & UATI Forwarding Tables and SC session lookup table

FIG. 7

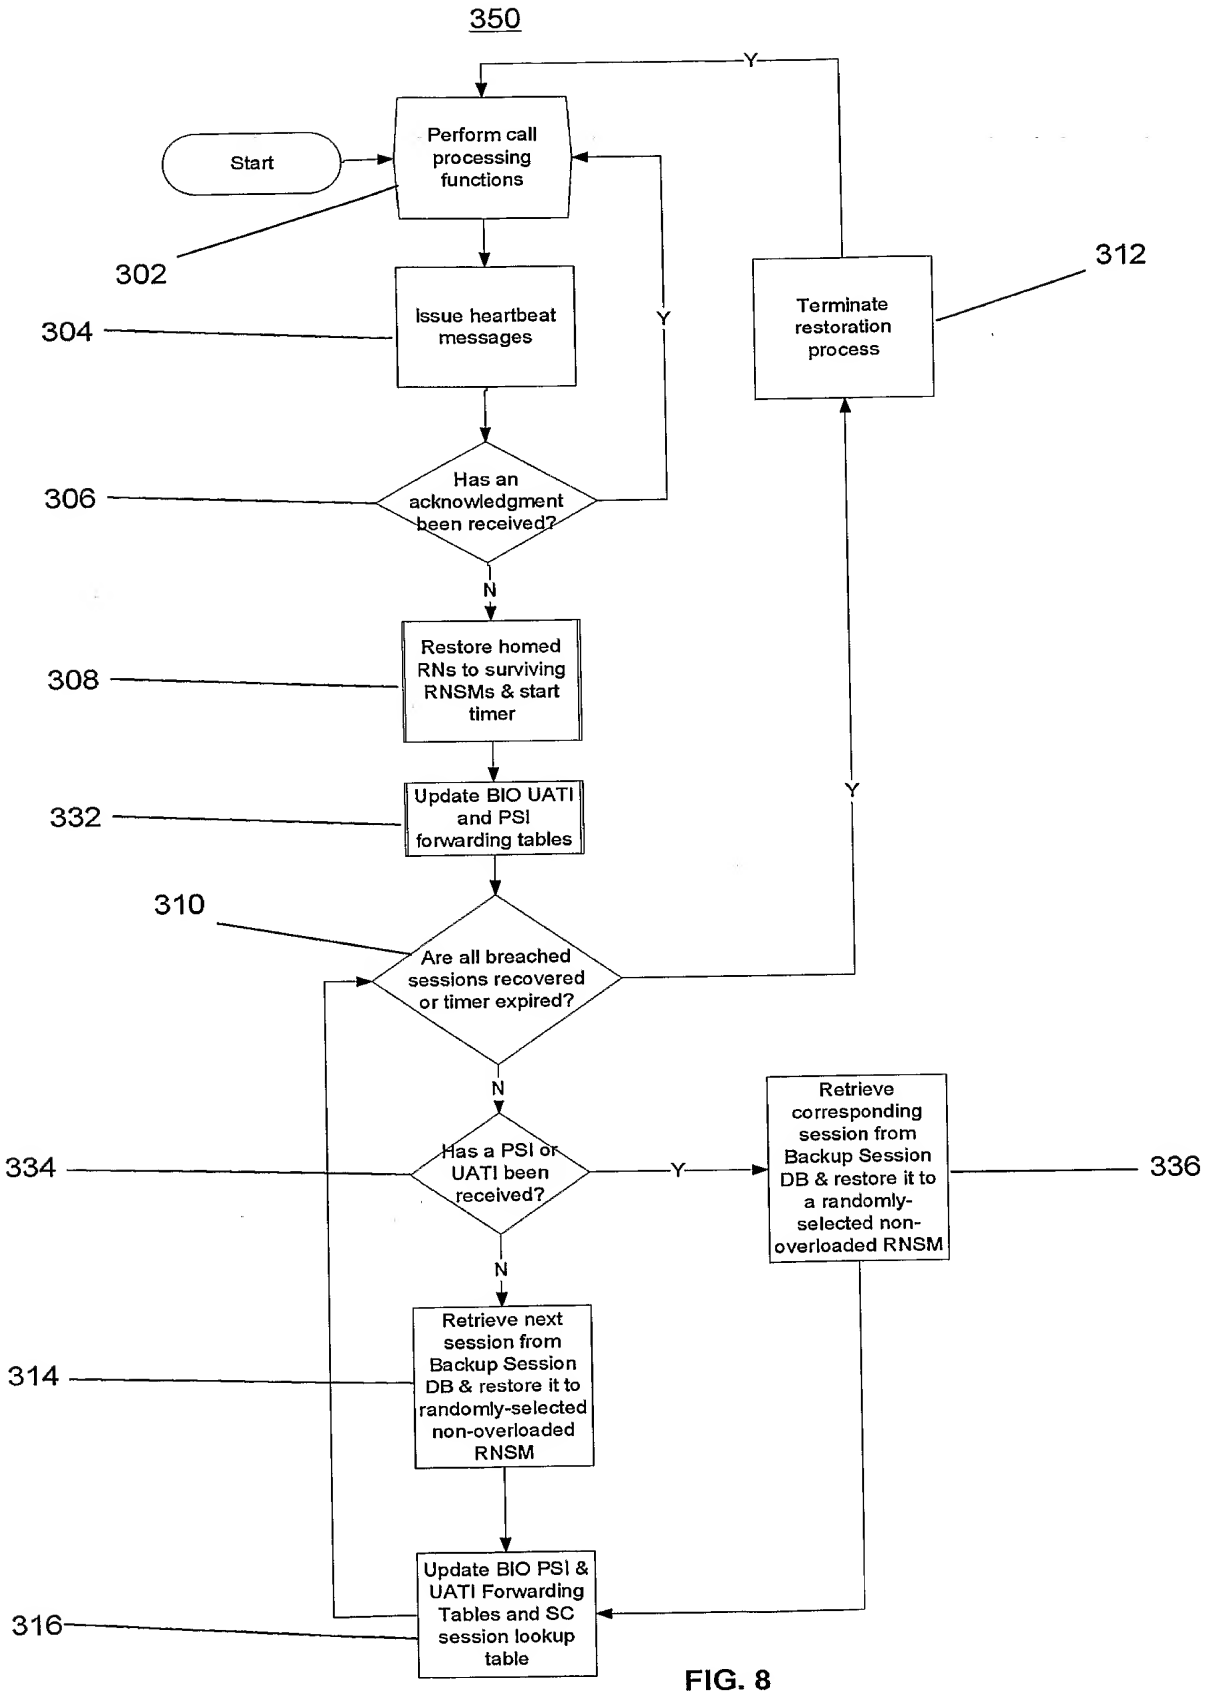
350
Start
Perform call processing functions
302
Issue heartbeat messages
304
Has an acknowledgment been received?
306
Y
N
Restore homed RNs to surviving RNSMs & start timer
308
Update BIO UATI and PSI forwarding tables
332
310
Are all breached sessions recovered or timer expired?
Y
N
Has a PSI or UATI been received?
334
Y
N
Retrieve corresponding session from Backup Session DB & restore it to a randomly-selected non-overloaded RNSM
336
Retrieve next session from Backup Session DB & restore it to randomly-selected non-overloaded RNSM
314
Update BIO PSI & UATI Forwarding Tables and SC session lookup table
316
Terminate restoration process
312

FIG. 8

	Session Index	1	2	3	4	...	N
254	UATI	UATI 1	UATI 2	UATI 3	UATI 4		UATI I
256	HwID	HwID 1	HwID 2	HwID 3	HwID 4		HWID I
258	PSI	PSI 1	PSI 2	PSI 3	PSI 4		PSI I
260	Control Cycle	32	23	12	10		CC I
262	Sector ID	Sector ID 1 [6]	Sector ID 2 [6]	Sector ID 3 [6]	Sector ID 4 [6]		Sector ID I [6]
264	Flags		Failed RNSM	Failed RNSM			
266	RNSM	RNSM 1	RNSM 2	RNSM 2	RNSM 1		RNSM I

FIG. 9

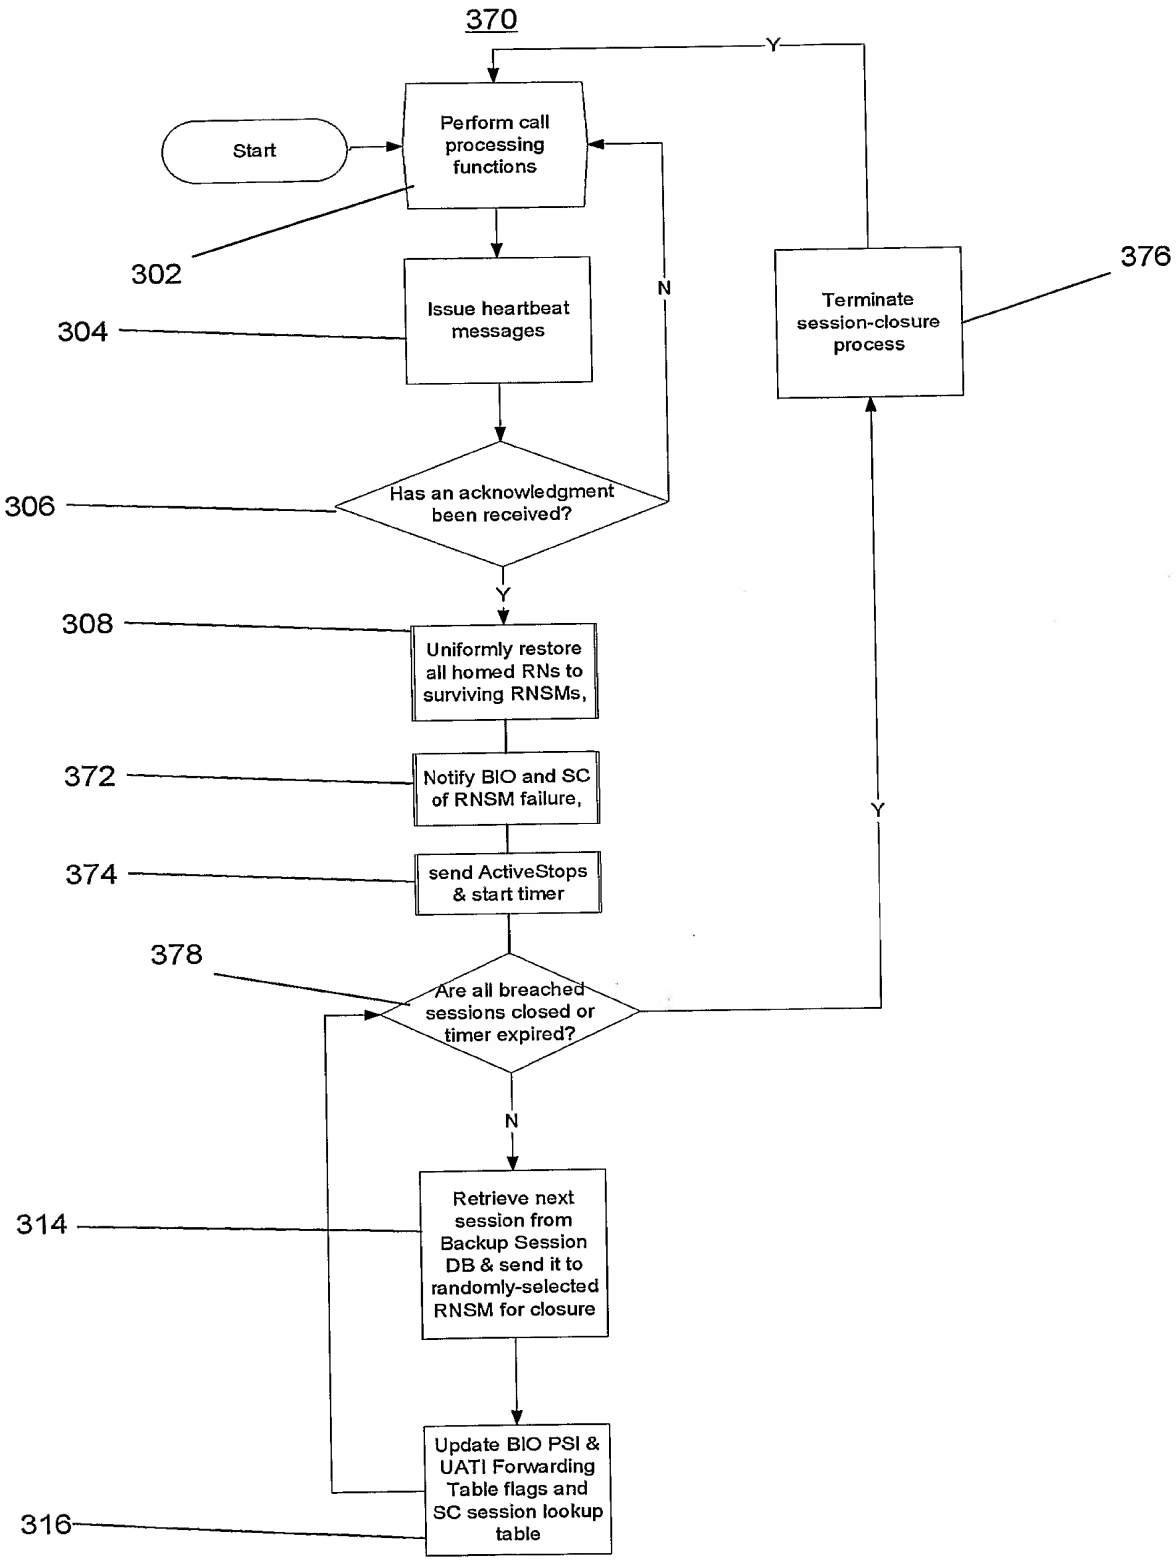
370
Start
Perform call processing functions
302
Issue heartbeat messages
304
Has an acknowledgment been received?
306
N
Y
Terminate session-closure process
376
Y
308
Uniformly restore all homed RNs to surviving RNSMs,
372
Notify BIO and SC of RNSM failure,
374
send ActiveStops & start timer
378
Are all breached sessions closed or timer expired?
N
314
Retrieve next session from Backup Session DB & send it to randomly-selected RNSM for closure
Update BIO PSI & UATI Forwarding Table flags and SC session lookup table
316

FIG. 10

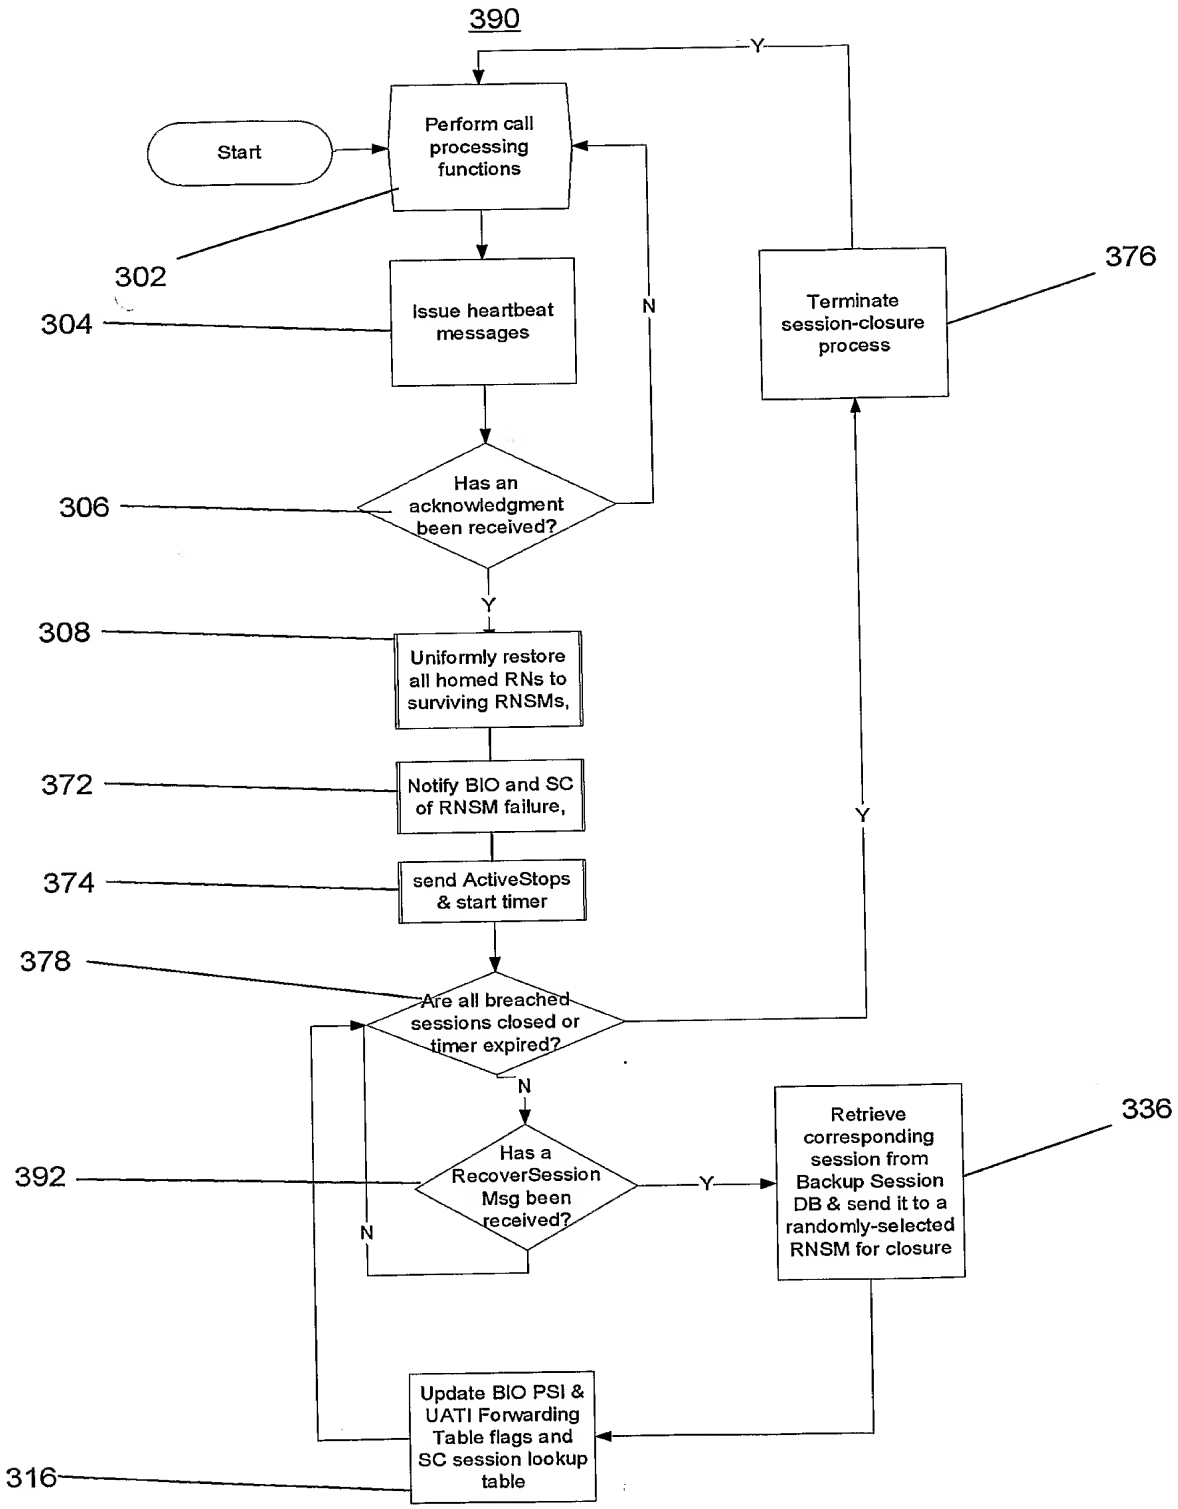
390
Start
Perform call processing functions
302
304
Issue heartbeat messages
306
Has an acknowledgment been received?
N
Y
308
Uniformly restore all homed RNs to surviving RNSMs,
372
Notify BIO and SC of RNSM failure,
374
send ActiveStops & start timer
378
Are all breached sessions closed or timer expired?
N
Y
392
Has a RecoverSession Msg been received?
N
Y
336
Retrieve corresponding session from Backup Session DB & send it to a randomly-selected RNSM for closure
376
Terminate session-closure process
Y
316
Update BIO PSI & UATI Forwarding Table flags and SC session lookup table

FIG. 11

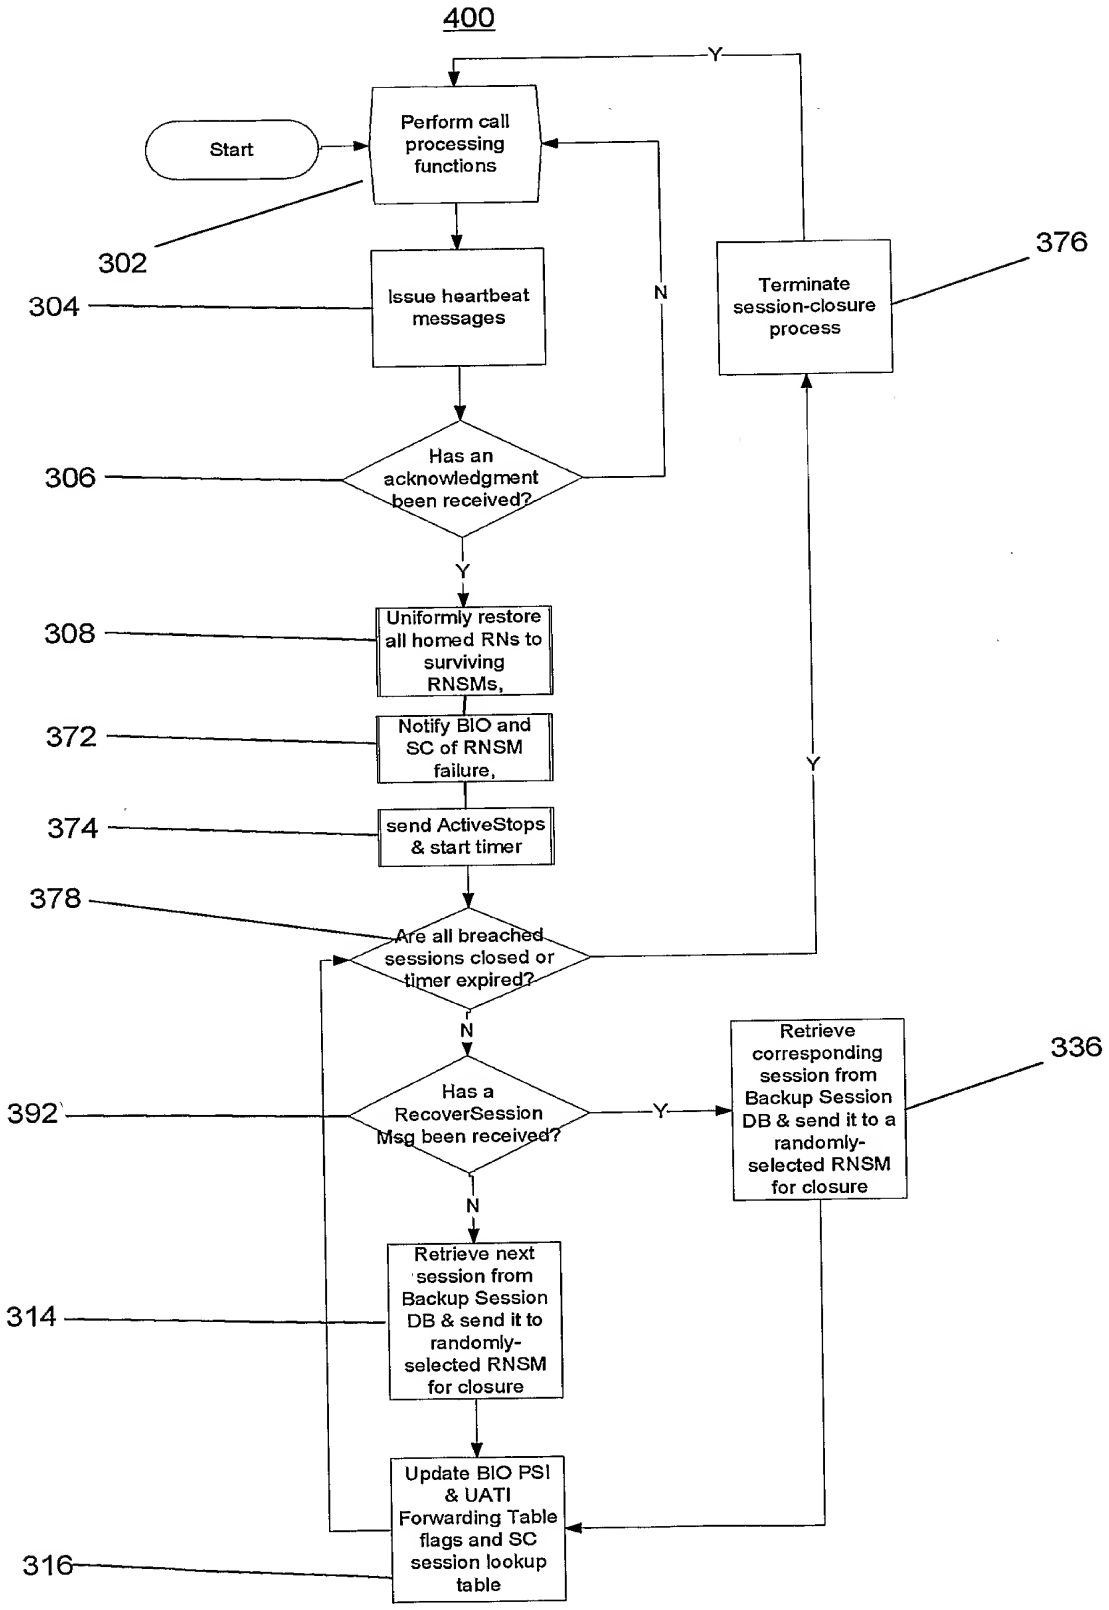
400
Start
Perform call processing functions
302
Issue heartbeat messages
304
Has an acknowledgment been received?
306
N
Y
Uniformly restore all homed RNs to surviving RNSMs,
308
Notify BIO and SC of RNSM failure,
372
send ActiveStops & start timer
374
Are all breached sessions closed or timer expired?
378
N
Y
Has a RecoverSession Msg been received?
392
Y
N
Retrieve corresponding session from Backup Session DB & send it to a randomly-selected RNSM for closure
336
Retrieve next session from Backup Session DB & send it to randomly-selected RNSM for closure
314
Update BIO PSI & UATI Forwarding Table flags and SC session lookup table
316
Terminate session-closure process
376
Y

FIG. 12

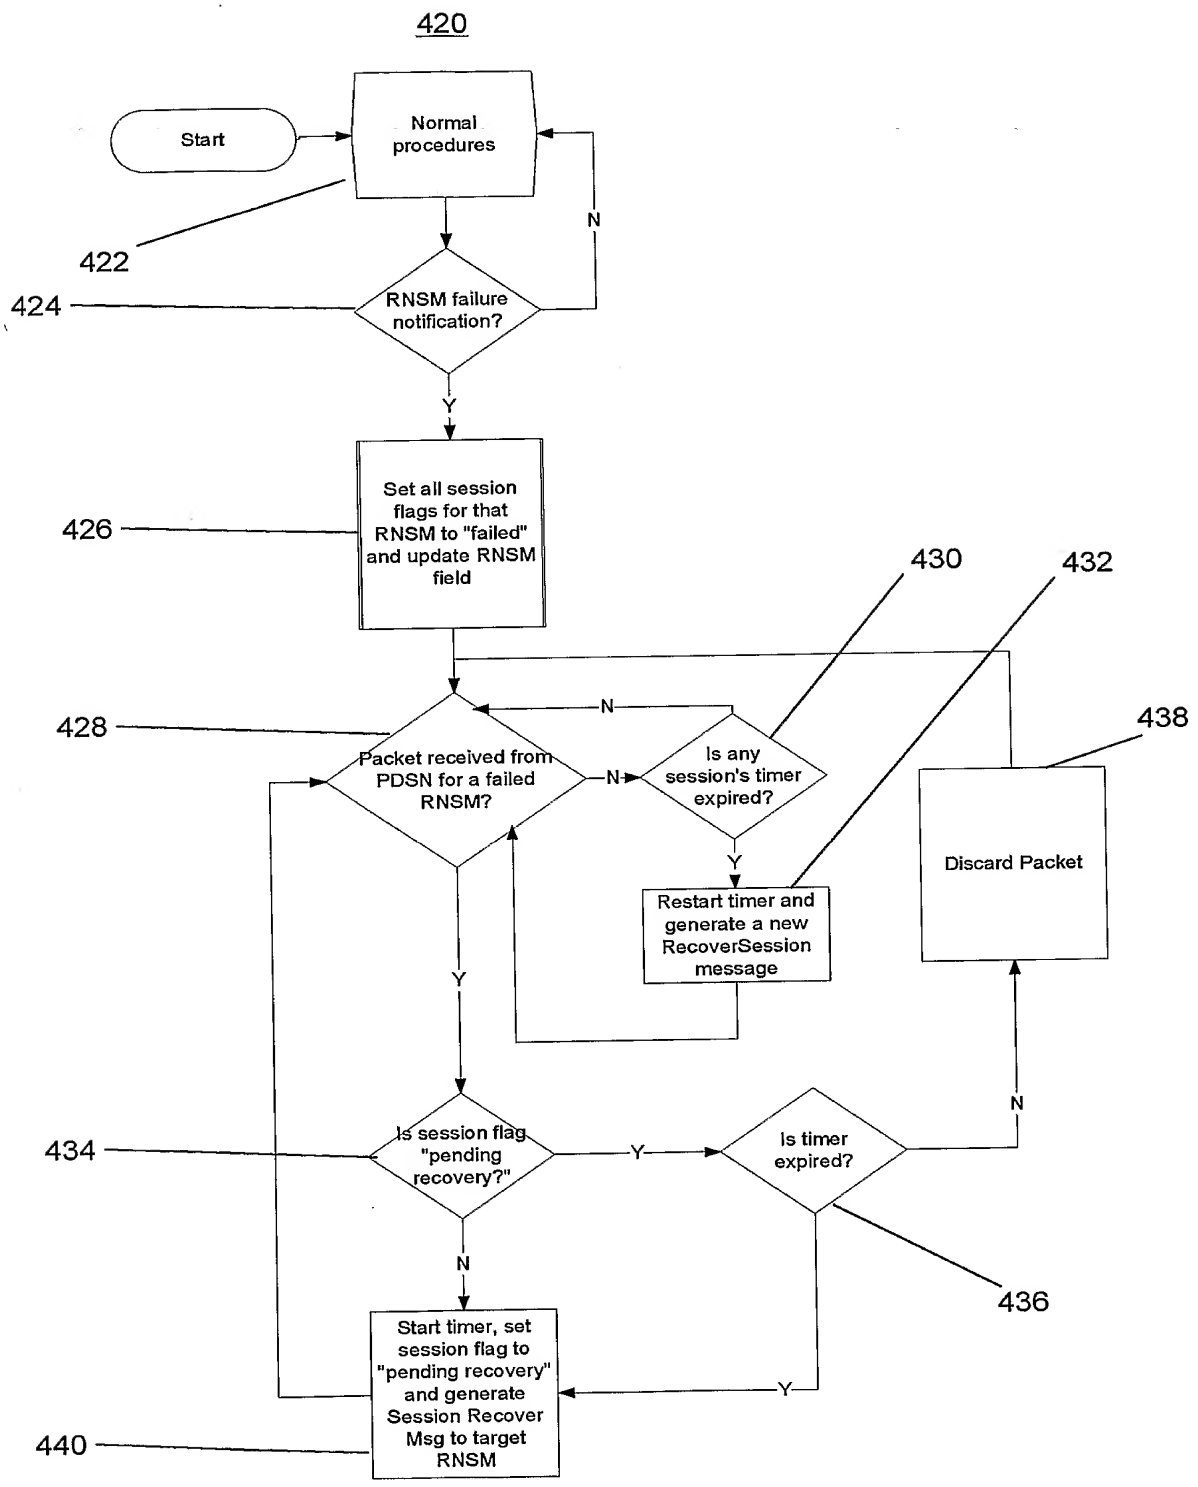
420
Start
Normal procedures
422
424
RNSM failure notification?
N
Y
426
Set all session flags for that RNSM to "failed" and update RNSM field
430
432
428
Packet received from PDSN for a failed RNSM?
N
Is any session's timer expired?
438
Discard Packet
Y
Restart timer and generate a new RecoverSession message
Y
N
434
Is session flag "pending recovery?"
Y
Is timer expired?
N
436
N
Start timer, set session flag to "pending recovery" and generate Session Recover Msg to target RNSM
Y
440

FIG. 13

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
19 April 2007 (19.04.2007)

PCT

(10) International Publication Number
WO 2007/044099 A3

(51) International Patent Classification:
G06F 15/16 (2006.01)

(21) International Application Number:
PCT/US2006/025018

(22) International Filing Date: 26 June 2006 (26.06.2006)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/166,893 24 June 2005 (24.06.2005) US

(63) Related by continuation (CON) or continuation-in-part (CIP) to earlier application:
US 11/166,893 (CON)
Filed on 24 June 2005 (24.06.2005)

(71) Applicant *(for all designated States except US)*: **AIRVANA, INC.** [US/US]; 19 Alpha Road, Chelmsford, Massachusetts 01824 (US).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **CHERIAN, Sanjay** [US/US]; 6 Maxwell Drive, Brookline, New Hampshire 03033 (US). **NG, Dennis** [US/US]; 126 Indian Meadow Drive, Northboro, Massachusetts 01532 (US). **BARABELL, Arthur J.** [US/US]; 11 Hayden Circle, Sudbury, Massachusetts 01776 (US). **RAMASWAMY, Suresh** [US/US]; 40 Old Stage Road, Chelmsford, Massachusetts 01824 (US). **GARG, Deepak** [IN/US]; 56 Stillwater Drive, Nashua, New Hampshire 03062 (US).

(74) Agents: **FEIGENBAUM, David L.** et al.; Fish & Richardson P.C., P.o. Box 1022, Minneapolis, Minnesota 55440 (US).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA,

[Continued on next page]

(54) Title: PRESERVING SESSIONS IN A WIRELESS NETWORK

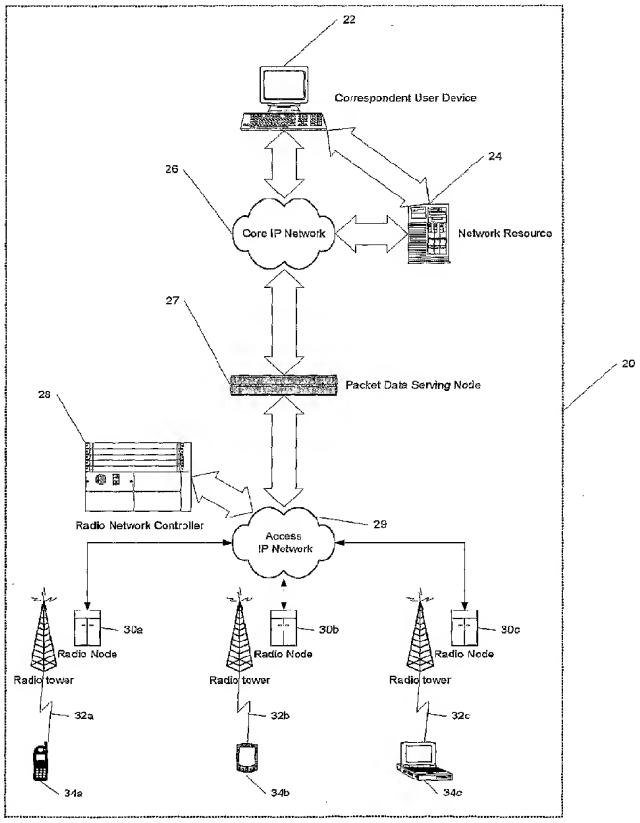

(57) Abstract: A radio network controller and methods for reestablishing sessions in a wireless network are described. At least a portion of session information associated with a first session is saved; and in response to detecting an unexpected degradation of the first session, reestablishment of the first session is triggered using the portion of the session information.

WO 2007/044099 A3

NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

— *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

— *with international search report*

(88) Date of publication of the international search report:
9 April 2009

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US 06/25018

A. CLASSIFICATION OF SUBJECT MATTER
IPC(8): G06F 15/16 (2007.01)
USPC: 709/227
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC(8): G06F 15/16 (2007.01); USPC: 709/227

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
USPC: 709/224, 227, 203; 370/338; 455/436; 711/118 (keyword limited -- see keywords below)

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

PubWEST(PGPB,USPT,EPAB,JPAB), SCHOLAR GOOGLE: mobile, drop, reestablish, session, wireless, maintain, open close, restoring, queue, degradation, breach

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X -- Y	US 2004/0038700 A1 (GIBBS) 26 February 2004 (26.02.2004), entire document; Abstract; para [0046]-[0051]; Fig 4-5; para [0034]; para [0048]; para [0051], [0054]; para [0047]; para [0053].	23-24 ---------- 1-22, 25-59
Y	US 2004/0008649 A1 (WYBENGA et al.) 15 January 2004 (15.01.2004), entire document; Abstract; para [0021]; para [0022]; para [0019]; para [0030]; para [0027]; para [0066]; para [0061]; para [0049]-[0051].	1-22, 25-41, 51-59
Y	US 6,542,481 B2 (FOORE et al.) 01 April 2003 (01.04.2003), entire document; Abstract; col 6, ln 49-51; col 7, ln 66-col 8, ln 45; Fig 4; col 3, ln 49-51; col 8, ln 1-17; col 8, ln 57-col 9, ln 3; col 9, ln 4-18; col 8, ln 13-16.	11-13, 15-16, 33-34, 42-50, 56-59
Y	US 2003/0031201 A1 (CHOI) 13 February 2003 (13.02.2003), entire document; Abstract; para [0019]-[0020], [0023], [0027]; para [0030].	8, 29, 43, 55

☐ Further documents are listed in the continuation of Box C. ☐

* Special categories of cited documents:
"A" document defining the general state of the art which is not considered to be of particular relevance
"E" earlier application or patent but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed
"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
06 August 2007 (06.08.2007)	29 JAN 2008
Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-3201	Authorized officer: Lee W. Young PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774

Form PCT/ISA/210 (second sheet) (April 2007)